

# Deriving Photometric Redshifts using Fuzzy Templates and Self-Organizing Maps. II. Comparing Sampling Techniques Using Mock Data

Joshua S. Speagle<sup>1,2\*</sup> and Daniel J. Eisenstein<sup>1</sup>

<sup>1</sup>*Harvard University Department of Astronomy, 60 Garden St., MS 46, Cambridge, MA 02138, USA*

<sup>2</sup>*Kavli IPMU (WPI), UTIAS, The University of Tokyo, Kashiwanoha 5-1-5, Kashiwa, Chiba, Japan*

Accepted XXX. Received YYY; in original form ZZZ

## ABSTRACT

In a companion paper (Speagle & Eisenstein 2015a), we proposed combining large numbers of “fuzzy archetypes” with Self-Organizing Maps (SOMs) to derive photometric redshifts in a data-driven way. In this paper, we investigate the performance of several sampling approaches that build on this general idea using a mock catalog designed to approximately simulate LSST (*ugrizY*) and *Euclid* (*YJH*) data from  $z = 0 - 6$  at fixed LSST  $Y = 24$  mag. We test eight different approaches: two brute-force methods, two Markov Chain Monte Carlo (MCMC)-based methods, two hierarchical sampling methods, and two “quick-search” methods based on quantities derived during the initial SOM training process. We find most methods perform reasonably well with small catastrophic outlier fractions and are able to robustly identify redshift probability distribution functions that are multi-modal and/or poorly constrained. Once these insecure objects are removed, the results are generally in good agreement with the strict accuracy requirements necessary to meet *Euclid* weak lensing goals for most redshifts above  $z \sim 0.8$ . These results demonstrate the utility of our data clustering-based approach and highlight its effectiveness to derive quick and accurate photo- $z$ ’s using large numbers of templates.

**Key words:** methods: statistical – techniques: photometric – galaxies: distances and redshifts

## 1 INTRODUCTION

Future large-scale surveys such as *Euclid* (Laureijs et al. 2011) and the Large Synoptic Survey Telescope (Ivezic et al. 2008) will rely on accurate redshifts ( $z$ ) measured to a significant portion of observed galaxies in order to measure the dark energy equation-of-state to target goals of a few percent accuracy using weak lensing (Albrecht et al. 2006; Bordoloi et al. 2012; Newman et al. 2015). Due to the challenge of obtaining spectroscopic redshifts (spec- $z$ ’s) to the majority galaxies in these samples, many of these studies will heavily rely on accurate photometric redshifts (photo- $z$ ’s) derived from galaxy broad- and/or narrow-band spectral energy distributions (SEDs).

Two main techniques are used to derive photo- $z$ ’s and the associated redshift probability distribution function (PDF)  $P(z|\mathbf{F}_{\text{obs}})$  for a given set of observed fluxes  $\mathbf{F}_{\text{obs}}$ . Template fitting approaches (e.g., Benítez 2000; Bolzonella

et al. 2000; Feldmann et al. 2006; Ilbert et al. 2006; Cool et al. 2013; Johnson et al. 2013; Speagle et al. 2015) rely on deriving a set of forward mappings and their associated likelihoods from a collection of templates and associated model parameters to observed color space. Machine learning approaches (e.g., Collister & Lahav 2004; Abazajian et al. 2009; Li & Yee 2008; Gerdes et al. 2010; Carrasco Kind & Brunner 2014a; Bonnett 2015; Hoyle 2015; Elliott et al. 2015; Sadeh et al. 2015; Almosallam et al. 2015), on the other hand, use training data to derive the best inverse mapping from observed color space to redshift.<sup>1</sup>

The majority of photo- $z$  developments in the last few years has been dominated by the introduction of numerous machine learning-based approaches. These are attractive relative to template fitting methods for several reasons:

<sup>1</sup> See Hildebrandt et al. (2010), Dahlen et al. (2013), and Sánchez et al. (2014) for recent overviews and comparisons, and Ménard et al. (2013) as well as related work for new applications of clustering-based redshifts (which will not be discussed here).

\* E-mail: jspeagle@cfa.harvard.edu

(i) *Single parameter estimation*: While template-fitting approaches are physically motivated, they involve fitting a number of extraneous parameters that must be marginalized over afterwards to derive  $P(z|\mathbf{F}_{\text{obs}})$ . Machine learning methods by design tend to focus on single-parameter estimation, making them somewhat more attractive.

(ii) *Computational speed*: Most template codes either resort to brute-force sampling (Ilbert et al. 2006; Cool et al. 2013) or Markov Chain Monte Carlo (MCMC) techniques (Johnson et al. 2013; Speagle et al. 2015) run over a large pre-generated grid of model photometry to derive  $P(z|\mathbf{F}_{\text{obs}})$ , both of which tend to take significantly more time than the majority of machine learning architectures.

(iii) *Empirically driven*: Most template fitting codes suffer from serious systematic uncertainties involved with modeling the impact of dust attenuation and emission line variation on an often limited set of templates (e.g., Coleman et al. 1980; Kinney et al. 1996; Polletta et al. 2007) along with the need for good relative photometric calibrations. Machine learning methods are not subject to these uncertainties.<sup>2</sup>

While machine learning approaches have been shown to be effective within the bounds of their training sets (Sánchez et al. 2014), attempts to extend their predictions beyond them have only been partially successful (Hoyle et al. 2015). Due to the limited regions of color space spanned by spec- $z$ 's training sets today (Masters et al. 2015) and the difficulty of obtaining reliable spectra for higher redshift sources, it is likely that template-based photo- $z$ 's will continue to serve a crucial role in upcoming imaging surveys.

This paper is the latest in a series that attempt to make template fitting-based approaches more competitive in order to keep pace with the rapid development in machine learning applications. In Speagle et al. (2015), we investigated the general features of typical pre-generated model “grids” and were able to develop a simulated annealing and ensemble MCMC (Goodman & Weare 2010; Foreman-Mackey et al. 2013) driven approach to exploring those grids more efficiently in an attempt to improve computational speed and efficiency. However, this did not affect the way that these grids of model photometry were created or improve on the modeling assumptions involved.

In an attempt to fix this glaring drawback, we introduced a new framework in Speagle & Eisenstein (2015a) (henceforth Paper I) that allows an algorithm to explore a large number of “fuzzy archetypes” using clustering methods such as Self Organizing Maps (SOMs). This approach significantly reduces systematics from dust and emission-line modeling and allows for an adaptive exploration of the space in a computationally efficient manner.

In this paper, we investigate the performance of several extensions of this new hybrid approach against a mock LSST and *Euclid* catalog. As the relevant likelihood surface is challenging to search, with many potential local minima that are often widely-separated (Speagle et al. 2015), we limit our focus in this paper to the case the same model parameters has been used to generate both the templates

and observed galaxies. This enables us to better investigate the relative performance of each algorithm and avoid issues relating to template-mismatch systematics.

The paper is organized as follows. In §2, we briefly summarize the basics of template-fitting approaches and the framework outlined in Paper I. In §3, we outline the eight basic methods we use to derive photo- $z$ 's using our fuzzy templates and SOMs. In §4, we describe our mock catalog and examine the performance of each of our individual approaches. We conclude in §5.

Throughout this work, boldface font ( $\mathbf{x}$ ,  $\Theta$ ) is used to represent vectors/matrices while normal font ( $x$ ,  $\theta$ ) is used for singular variables. All observations are generated from 9 bands of data (LSST *ugrizY* and *Euclid YJH*) with error properties based on  $5\sigma$  imaging depths of 26.5 mag (24.5 in *u*) for LSST and 24 mag for *Euclid* along with an assumed calibration uncertainty of 0.01 mag ( $\approx 1\%$ ).

## 2 FORMALISM

### 2.1 Fuzzy Templates

Model photometric fluxes  $\mathbf{F}_{\text{model}}$  for a given series of templates are generated via

$$\mathbf{F}_{\text{model}} = \frac{\int_{\nu_z} S_{\nu,\text{model}}(\nu) R_{\text{model}}(\nu) \mathbf{T}(\nu) \nu^{-1} d\nu}{\int_{\nu_z} \mathbf{T}(\nu) \nu^{-1} d\nu}, \quad (1)$$

where the model template

$$S_{\nu,\text{model}}(\nu) = \sum_{\text{gal}} c_{\text{gal}} \times S_{\nu,\text{gal}}(\nu) + \sum_{\text{lines}} \text{EW}_{\text{lines}} \times S_{\nu,\text{lines}}(\nu) \quad (2)$$

is a linear combination of “basis” galaxy (see, e.g., Coleman et al. 1980; Kinney et al. 1996; Polletta et al. 2007) and emission line templates (e.g., Kennicutt & Evans 2012) scaled by their equivalent widths (EWs), the dust (i.e. reddening) model

$$-2.5 \log(R_{\text{model}}(\nu)) = \sum_{\text{dust}} E(B-V)_{\text{dust}} k'_{\text{dust}}(\nu) + A_{\text{IGM}}(\nu, z) \quad (3)$$

is the sum of a linear combination of dust templates for a given  $E(B-V)$  that are applied as a uniform screen, with the rest-frame dust screen  $k'_{\text{dust}}(\nu) = k_{\text{dust}} + c_{b,\text{dust}} k_{\text{bump}}$  composed of a combination of an underlying dust continuum and the 2175 Å bump (Fitzpatrick & Massa 2007) and the IGM reddening template given by the parameterization outlined in Madau (1995). This combined template is convolved with the corresponding set of filter curves  $\mathbf{T}(\nu)$  and over the redshifted set of frequencies  $\nu_z$  probed by the respective filters.

As many template-fitting approaches only explore a single basis galaxy and dust template at a given time, we henceforth ignore linear combinations of those respective components. In addition, for convenience we will re-parameterize our full set of parameters into two vectors:  $\theta = \{z, \text{gal}, \text{dust}\}$ , which contains the parameters surrounding the baseline archetype, and  $\phi = \{E(B-V), c'_b \equiv c_b \times E(B-V), \Delta \text{EW}\}$ , which contains the associated “fuzzy” parameters.

For computational reasons (see Speagle et al. 2015),

<sup>2</sup> If photo- $z$ 's are used to fit stellar population synthesis models at a later period, however, then good relative photometric calibrations still are required to avoid introducing biases.

most codes opt to use large *pre-generated grids* of photometry from combinations of parameters sampled at a given granularity in each dimension during the fitting process. However, most of these grids are constructed using only a small collection of galaxy ( $\lesssim 30$ ) and dust ( $\lesssim 5$ ) templates taken from low- $z$  observations and then subsequently applied to the high- $z$  universe with large modifications.

Instead of starting with a small number of templates in well-understood regions of color space and trying to simulate less-understood regions of color space, we instead attempt to sample the color-space manifold occupied by galaxies directly using a much larger set of galaxy spectral “archetypes”. In order to allow for possible variation away from our discrete set of starting archetypes, we generate a small level of intrinsic “fuzziness” using sets of individualized, independent priors on dust attenuation and emission line strengths that allows for small deviations away from each baseline archetype. This transforms an individual galaxy’s position in color space into a multidimensional PDF centered at its original location. By keeping the perturbations relatively small, we remain in the regime where the uniform dust screen approximation is likely to remain valid while also allowing emission line strengths to vary independently of one another.

As Paper I, we use the set of 129 high-quality UV–IR spectra from Brown et al. (2014) to serve as our baseline set of archetypes and measure the corresponding EWs of  $\{H\alpha+N[III], H\beta, H\gamma, O[III], \text{ and the } O[III] \text{ doublet}\}$  to derive appropriate emission line templates. These are combined with a series of 8 dust curves of the form  $k(x) = a_x(\beta)(x)^{-\beta} + b_x(\beta) + c_b k_{\text{bump}}(x|x_0, \gamma)$ , where  $\beta = \{0.3, 0.45, 0.6, 0.75, 0.9, 1.05, 1.2, 1.35\}$ ,  $a_x$  and  $b_x$  are chosen such that each dust curve is normalized to  $A(V) = 3.0$  and  $A(2.5 \mu\text{m}) = 0$ , and the mean amplitude ( $c_b$ ), position ( $x_0$ ), and width ( $\gamma$ ) of the 2175 Å bump are taken from the mean fits presented in Fitzpatrick & Massa (2007). We then choose  $P(\phi) = P[\Delta E(B-V)] \times P[\Delta c'_b] \times P[\Delta \text{EW}_{\text{line}}]$  such that

$$\begin{aligned} P[\Delta E(B-V)] &= \mathcal{N}[0, \Delta A(1500 \text{ Å}) = 0.05 \text{ mag}], \\ P[\Delta c'_b] &= \mathcal{N}[0, 1.5] \times P[\Delta E(B-V)], \text{ and} \\ P[\Delta \text{EW}_{\text{line}}] &= \mathcal{N}[0, 0.2 \times \max(\text{EW}, 0.005 \lambda_{\text{line}})], \end{aligned} \quad (4)$$

where  $\mathcal{N}(\mu, \sigma)$  indicates a normal (i.e. Gaussian) distribution with mean  $\mu$  and standard deviation  $\sigma$ . These form the underlying set of priors for our fuzzy archetypes.

Our decision to fix  $P[\Delta E(B-V)]$  to  $A(1500)$  limits the amount of dust extinction in the UV to the approximately linear regime ( $\sim 5\%$   $1\sigma$  variation) across the entire wavelength range for each dust curve. Likewise, our choice of prior for  $\Delta c'_b$  allows fluctuations to 0 to take place at the  $2\sigma$  level, which enables us to fit dust curves without the impact of the 2175 Å bump while still disfavoring possible non-physical solutions. Finally, the 20% variation ( $\sim 0.1$  dex) in EW allowed by  $P(\Delta \text{EW}_{\text{line}})$  mimics the observed scatter among emission line scaling relations (Kennicutt 1998; Kennicutt & Evans 2012), while the EW “floor” leaves open the possibility for small variation at the  $\lesssim 1-2\%$  level.

## 2.2 Self-Organizing Maps

The Self-Organizing Map (SOM; Kohonen 1982; Kohonen 2001) is an unsupervised machine learning algorithm that

projects high-dimensional data onto a lower-dimensional space using competitive training of a large set of “cells” in a way that preserves general topological features and correlations present in the higher-dimensional data. Each cell in the grid is assigned a position  $\mathbf{x}_{\text{som}}^{\text{cell}}$  and contains a cell model  $\mathbf{F}_{\text{som}}^{\text{cell}}(t)$  that is the same size as a vector from the training data ( $N_{\text{filter}}$ ) and initialized (most often randomly) at iteration  $t = 0$ . Training then proceeds as follows:

- (i) Draw (with replacement) a random object  $\mathbf{F}_{\text{obs}}^i$  from the input dataset.
- (ii) Compute  $\chi^2(\mathbf{F}_{\text{obs}}^i, \mathbf{F}_{\text{som}}^{\text{cell}}(t), s_{\text{cell}}(t))$  for every cell on the map.
- (iii) Select the best-matching (i.e. lowest  $\chi^2$ ) cell  $\mathbf{x}_{\text{som}}^{\text{best}}$  from the corresponding set of  $\chi^2$ 's.
- (iv) Update the cell models within the map according to an evolving *learning rate*  $\mathcal{A}(t)$  and *neighborhood function* (i.e. kernel density)  $\mathcal{H}(\mathbf{x}_{\text{som}}^{\text{best}}, \mathbf{x}_{\text{som}}^{\text{cell}}|t)$  such that

$$\mathbf{F}_{\text{som}}^{\text{cell}}(t+1) = s_{\text{cell}}(t)\mathbf{F}_{\text{som}}^{\text{cell}}(t) + \mathcal{A}(t)\mathcal{H}(t) [\mathbf{F}_{\text{obs}}^i - s_{\text{cell}}(t)\mathbf{F}_{\text{som}}^{\text{cell}}]. \quad (5)$$

- (v) Increment  $t$  and repeat steps (i) through (iv) while  $t < N_{\text{iter}}$ .

After training, objects are “sorted” onto the map by repeating steps (ii) and (iii) for every object in the input dataset and assigning them to the best-matching cell.

Following Paper I, we set

$$\begin{aligned} \mathcal{A}(t) &= a_0 \left( \frac{a_1}{a_0} \right)^{t/N_{\text{iter}}}, \\ \mathcal{H}(t|\mathbf{x}_{\text{som}}^{\text{best}}, \mathbf{x}_{\text{som}}^{\text{cell}}) &= \exp \left( - \frac{\|\mathbf{x}_{\text{som}}^{\text{best}} - \mathbf{x}_{\text{som}}^{\text{cell}}\|^2}{\sigma^2(t)} \right), \text{ and} \\ \sigma(t) &= \sigma_0 \left( \frac{1}{\sigma_0} \right)^{t/N_{\text{iter}}}, \end{aligned} \quad (6)$$

where  $\mathcal{A}(t)$  and  $\mathcal{H}(t)$  are now functions of the hyperparameters  $a_0$ ,  $a_1$ , and  $\sigma_0$ .

We use the SOM to organize both rest-frame and observed-frame colors using “model catalogs” generated from numerous Monte Carlo realizations of our fuzzy templates and a given  $P(\theta)$ . These are generated as follows:

- (i) *Rest-frame SOM* (2-D): We construct a model catalog of 9,675 objects using 75 Monte Carlo realizations from  $P(\phi)$  and uniform sampling of  $k_{\text{dust}}$  for each fuzzy archetype, which are assigned 1% errors to in order to ensure the SOM training proceeds in logarithmic rather than linear space. We construct a two-dimensional  $30 \times 30$  SOM and initialize the cell models randomly on the uniform interval  $[0, 1]$ . We set  $[a_0, a_1, \sigma_0] = [1, 0.5, 30]$  and allow the map to train for  $N_{\text{iter}} = 10,000$  iterations. The corresponding set of mappings from the model catalog onto the SOM are then recorded.

- (ii) *Observed-frame SOM* (3-D): We construct a model catalog of 3,228,225 objects using 25 Monte Carlo realizations from  $P(\phi)$  with uniform sampling from  $k_{\text{dust}}$  for each fuzzy archetype sampled on every point of an input  $z = 0-10$  ( $\Delta z = 0.01$ ) redshift grid. Each object is assigned 1% errors as well as a baseline error floor of 0.01, which are added in quadrature. We construct a three-dimensional  $30 \times 30 \times 30$  SOM and initialize the cell models according to pre-defined color gradients in  $u-r$ ,  $r-z$ , and  $z-H$ .

We set  $[a_0, a_1, \sigma_0] = [0.5, 0.2, 10]$  and allow the map to train for  $N_{\text{iter}} = 1,000,000$  iterations. The corresponding set of mappings from the model catalog onto the SOM are then recorded.

### 3 METHODS

To test whether our use of (linearized) fuzzy templates and Self-Organizing Maps (SOMs) can be used to derive photometry, we decide to compare the following eight different methods:

(i) **BruteForce**: A standard brute-force approach that involves comparing against all 3,228,225 models used to create the 3-D observed-frame SOM.

(ii) **BruteForce\_LinearFuzzy**: A modified brute-force approach that explores all  $N_z \times N_{\text{gal}} \times N_{\text{dust}} = 1,033,032$  possible combinations of our non-linear parameters  $\theta$ , where at each step we optimize over all linear parameters  $\phi$ .

(iii) **SOM\_MCMC\_RestFrame**: An MCMC-based approach that explores a 2-D SOM constructed with rest-frame realizations of our fuzzy templates along with changes in photometry as a function of redshift.

(iv) **SOM\_MCMC\_ObservedFrame**: An MCMC-based approach that explores a 3-D SOM constructed from *observed-frame* realizations (i.e. redshift already included) of our fuzzy templates.

(v) **SOM\_Hierarchical\_MonteCarlo**: A hierarchical sampling approach that uses a fixed number of randomly sampled objects drawn from each SOM cell.

(vi) **SOM\_Hierarchical\_ImportanceSampling**: A hierarchical sampling approach that draws objects randomly from SOM cells based on an adaptive SOM importance density.

(vii) **SOM\_CellModel\_LimitedSum**: A limited reconstruction using all models contained within a small subset of SOM cells selected based on their SOM cell models.

(viii) **SOM\_CellModel\_Average**: A weighted redshift average across the SOM based on the SOM cell models.

We discuss each of these in turn.

#### 3.1 BruteForce (gold standard)

Deriving  $P(z|\mathbf{F}_{\text{obs}})$  for any individual object  $\mathbf{F}_{\text{obs}}$  in a brute-force approach involves computing the likelihood  $P(\theta, \phi|\mathbf{F}_{\text{obs}})$  to the entire collection  $\mathbf{F}_{\text{model}}(\theta, \phi)$ 's. Each  $\mathbf{F}_{\text{model}}$  has a corresponding redshift attached to it, and the final  $P(z|\mathbf{F}_{\text{obs}})$  is constructed by marginalizing over all the other input parameters such that

$$P(z|\mathbf{F}_{\text{obs}}) \propto \sum_{\text{dust}} \sum_{\text{gal}} \sum_{\phi} w(\theta, \phi) P(\theta, \phi|\mathbf{F}_{\text{obs}}), \quad (7)$$

where  $w(\theta, \phi)$  is an associated model weight (from, e.g., priors on  $\theta$  and  $\phi$ ) and  $P(\theta, \phi|\mathbf{F}_{\text{obs}})$  is the computed likelihood between the data and the model. For the purposes of this paper, we take our collection of model photometry to be the 3,228,225 objects included in our observed-frame model catalog (see §2.2) and assign all models equal weight such that  $w(\theta, \phi) = 1$ .

#### 3.2 BruteForce\_LinearFuzzy

In order to analytically marginalize over the “fuzziness” in our fuzzy templates, we first need to render them in a more statistically-friendly and computationally-convenient form. As outlined in Paper I, we can accomplish this by approximating the changes in the underlying archetype as being linear in  $\phi$  and recasting the fitting procedure as a (conditional) linear algebra problem that can be solved using iterative least squares. We detail the full derivation in Appendix A and give a brief outline below.

Taylor expanding around the baseline archetype and ignoring all terms higher than first order,  $\mathbf{F}_{\text{model}}(z)$  can be written as

$$\begin{aligned} s\mathbf{F}_{\text{model}}(z) = & s\mathbf{F}_{\text{gal}}(z) + \Delta E(B-V) s\mathbf{F}_{\text{dust}}(z) \\ & + \Delta c'_b s\mathbf{F}_{\text{bump}}(z) + \sum_{\text{lines}} \Delta E W_{\text{lines}}^{\text{gal}} s\mathbf{F}_{\text{lines}}^{\text{gal}}(z), \end{aligned} \quad (8)$$

where  $\mathbf{F}_{\text{gal}}(z)$  is the flux of the baseline galaxy archetype,  $\mathbf{F}_{\text{dust}}(z)$  is the “dust flux”,  $\mathbf{F}_{\text{bump}}(z)$  is the “2175 Å bump flux”, and  $\mathbf{F}_{\text{lines}}^{\text{gal}}(z)$  is the corresponding emission line flux for a given archetype and emission line.

Assuming that our priors are Gaussian, the corresponding log-likelihood is

$$-2 \log P(\theta, \phi, s|\mathbf{F}_{\text{obs}}) \equiv \sum_i \sigma_i^{-2} [\Delta F_i(\theta, \phi, s)]^2 + \sum_j \left( \frac{\phi_j}{\sigma_{\phi_j}(\theta)} \right)^2 \quad (9)$$

where  $\sigma^2$  is the observed variances (we take all model errors to be 0),  $\Delta \mathbf{F}(\theta, \phi, s) = \mathbf{F}_{\text{obs}} - s\mathbf{F}_{\text{model}}(\theta, \phi)$  are the associated flux residuals,  $\sigma_{\phi}(\theta)$  are the standard deviations of the corresponding Gaussian priors for a given  $\theta$ , the sum over  $i$  is taken over all observed bands, and the sum over  $j$  is taken over all relevant nuisance parameters. This allows us to re-write our expression for  $\mathbf{F}_{\text{model}}$  as

$$sF_{\text{model},i}(\theta, \phi) = sF_{\text{gal},i}(\theta) + \sum_j s\mathbf{X}_{ij}(\theta)\phi_j, \quad (10)$$

where  $\mathbf{X}(\theta)$  is an  $N_{\text{filt}} \times (2 + N_{\text{lines}})$  matrix of pre-computed coefficients for a given  $\theta$ .

For fixed  $\theta$  and  $\phi$ , we can marginalize over  $s$  to minimize  $\chi_{\text{mod}}^2(s|\theta, \phi)$  via

$$s = \sum_i \sigma_i^{-2} F_{\text{obs},i} F_{\text{model},i} / \sum_i \sigma_i^{-2} F_{\text{model},i} F_{\text{model},i}, \quad (11)$$

while for fixed  $s$  and  $\theta$ , we can marginalize over  $\phi$  by solving

$$\left( \mathbf{X}(\theta)^T \mathbf{W}_{\text{obs}} \right) \Delta \mathbf{F}_{\text{gal}}(\theta) = \left( \mathbf{X}(\theta)^T \mathbf{W}_{\text{obs}} \mathbf{X}(\theta) + \mathbf{W}_{\phi}(\theta) \right) \phi, \quad (12)$$

where  $\Delta \mathbf{F}_{\text{gal}}(\theta) = \mathbf{F}_{\text{obs}} - s\mathbf{F}_{\text{gal}}(\theta)$  is the baseline galaxy flux residual,  $\mathbf{W}_{\text{obs}} = \text{diag}(\dots, \sigma_i^{-2}, \dots)$  is the associated observational weight matrix,  $\mathbf{W}_{\phi} = \text{diag}(\dots, \sigma_{\phi_j}^{-2}, \dots)$  is the prior weight matrix, and  $T$  is the transpose operator.

This gives us an simple iterative scheme for computing  $P_{\text{max}}(\theta|\mathbf{F}_{\text{obs}})$  with respect to  $s$  and  $\phi$  for a given choice of  $\theta$ . Using this approach, we execute a brute-force search over all  $N_z \times N_{\text{gal}} \times N_{\text{dust}} = 1,033,032$  possible combinations of



our non-linear parameters  $\theta$ . This gives us

$$P(z|\mathbf{F}_{\text{obs}}) \propto \sum_{\text{dust}} \sum_{\text{gal}} P_{\text{max}}(\theta|\mathbf{F}_{\text{obs}}). \quad (13)$$

### 3.3 SOM\_MCMC\_RestFrame

In Paper I, we introduced a method to derive estimates to  $P(z|\mathbf{F}_{\text{obs}})$  in a more exact sense using Markov Chain Monte Carlo (MCMC) methods running “over” the SOM whose performance only depends on the ability of the SOM to effectively cluster the space in a topologically smooth fashion. We summarize our approach here and direct the reader to Paper I for additional details.

In brief, we implement a hybrid version of MCMC sampling that transitions to new regions based on comparisons between individual sets of model photometry but with chains that live “on top of” the SOM. Chains propose new models in a hierarchical fashion by first proposing a new cell position<sup>3</sup> and then selecting a random model within that cell. Each chain then transitions to the proposed cell with an associated transition probability that satisfies detailed balance.

After the chains have converged to the stationary target distribution (the “burn-in” phase), they are allowed to sample for a set number of iterations. The final  $P(z|\mathbf{F}_{\text{obs}})$  can then be reconstructed via

$$P(z|\mathbf{F}_{\text{obs}}) \approx \sum_{\text{chain}} w_{\text{chain}} P(z|\text{chain}, \mathbf{F}_{\text{obs}}). \quad (14)$$

where  $w_{\text{chain}} = P_{\text{max}}(\mathbf{F}_{\text{obs}}|\mathbf{F}_{\text{model}, \text{chain}})$  is a maximum-likelihood weight used to suppress chains that happen to converge to spurious minima and

$$P(z|\text{chain}, \mathbf{F}_{\text{obs}}) = \sum_{\text{dust}}^{\text{chain}} \sum_{\text{gal}}^{\text{chain}} \sum_{\phi}^{\text{chain}} w(\theta, \phi|\mathbf{F}_{\text{obs}}) \quad (15)$$

is the sum over the trials from each individual chain with  $w(\theta, \phi|\mathbf{F}_{\text{obs}})$  the associated weight of each trial. This is usually taken to be  $w = 1$  if the trial was accepted and  $w = 0$  if the trial was rejected, although other weighting schemes are possible (Bernton et al. 2015).

We implement this sampling scheme on the 2-D rest-frame SOM with where new cells are proposed according to a multivariate Gaussian with  $\Sigma = \text{diag}[(\sigma_{\text{som}}^{\text{dim.1}})^2, (\sigma_{\text{som}}^{\text{dim.1}})^2]$  where  $\sigma_{\text{som}}^{\text{dim.1}} = \sigma_{\text{som}}^{\text{dim.2}} = 2$ . Once a specific archetype, dust template, and  $\phi$  have been chosen, we generate the associated redshifted photometry based on a separate Gaussian proposal with  $\sigma_z = 2.5\%$  of the corresponding redshift grid.

To improve performance and reduce the dependence on our initial choice of parameters, we include a short learning step during every MCMC run after the burn-in phase is complete where the corresponding size of each component of the proposal ( $\{\sigma_{\text{som}}^{\text{dim.1}}\}_i, \sigma_z$ ) is adjusted to better approximate the underlying shape of the PDF. We set a fixed number of iterations for burn-in, learning, and sampling of 5000, 1000, and 5000 trials, respectively, and sample each object using 32 chains run in parallel.

<sup>3</sup> If a proposed cell has no objects, the process is repeated from the proposed cell until a valid cell is reached.

### 3.4 SOM\_MCMC\_ObservedFrame

In §3.3, we used a 2-D SOM with 900 cells to sort 9675 rest-frame relations of our fuzzy archetypes that were then allowed to vary as a function of redshift. However, we can instead attempt to reorganize the entire fuzzy archetype-redshift *observed-frame* color space using a single SOM. We use this approach to construct an 3-D observed-frame SOM composed of 27,000 cells based on  $\sim 3 \times 10^6$  individual realizations of our model photometry from  $z = 0 - 10$ .

To explore this observed-frame SOM, we again implement the MCMC sampling scheme described above on the 3-D observed-frame SOM. New cells are proposed according to the multivariate Gaussian distribution with  $\Sigma = \text{diag}[(\sigma_{\text{som}}^{\text{dim.1}})^2, (\sigma_{\text{som}}^{\text{dim.2}})^2, (\sigma_{\text{som}}^{\text{dim.3}})^2]$  where  $\sigma_{\text{som}}^{\text{dim.1}} = \sigma_{\text{som}}^{\text{dim.2}} = \sigma_{\text{som}}^{\text{dim.3}} = 2$ . As before, we set a fixed number of iterations for burn-in, learning, and sampling of 5000, 1000, and 5000 trials, respectively, and sample each object using 32 chains run in parallel.

### 3.5 SOM\_Hierarchical\_MonteCarlo

Rather than taking advantage of the topologically smooth nature of the SOM to “explore” its corresponding reduced-dimensional manifold, we can instead treat the SOM as simply a way to partition the model space into a series of clusters (i.e. cells). Rather than using the relationship between these clusters (i.e. the “graph” they form) to define some sampling algorithm, we can then treat each cluster as an independent entity. Thus, instead of sampling from a local neighborhood function in SOM space to determine what quasi-contiguous regions of the SOM contain the majority of the likelihood, we instead wish to determine which isolated clusters contain significant amounts of likelihood to merit further exploration and/or sampling.

In other words, we want to break down the problem into a *hierarchical sampling* approach, where we first consider which clusters (stored in some arbitrary, hierarchical structure) to explore before considering the individual models themselves. In the case of the SOM, where the clusters are organized along a single manifold, this is a simple two-tiered problem involving the individual cells and the objects contained within them.

The process of determining how likely a particular  $\mathbf{F}_{\text{obs}}$  (along with associated errors  $\sigma_{\text{obs}}$ ) is matched to a particular cluster is an open problem. As clusters can have very different distributions of galaxy colors and occupation rates depending of its location in color space (Paper I), the use of simple summary statistics will often introduce biases depending on the properties data in question (see §4).

Alternately, we can directly sample from the PDF spanned by each cluster using Monte Carlo techniques. Unlike the case of summary statistics, this approach is guaranteed to give a good approximation to the likelihood contained within each cluster when the number of samples  $N_{\text{MC}} \sim N(\mathbf{x}_{\text{som}}^{\text{cell}})$ .

In addition to giving us information on the likelihood contained within each cluster, such an approach naturally introduces a *stratified Monte Carlo* sampling scheme to uncover the underlying PDF, where random model photometry drawn from partitioned regions of color space can be used to stochastically reconstruct the color-redshift relation probed

by  $\mathbf{F}_{\text{obs}}$  and established by the individual SOM cells. In the case outlined above where  $N_{\text{MC}} \sim N(\mathbf{x}_{\text{som}}^{\text{cell}})$  for each cell, this approach simply reduces to a Monte Carlo version of a brute-force approach.

However, the power of this stratified Monte Carlo approach is especially apparent in the sparse-sampling regime where only a small amount of objects is drawn from each cluster. While such an approach might not be effective in determining the precise likelihood contained within individual cluster, it does provide a fair statistical representation of a sample galaxy’s distribution over the entire region of color space spanned by the collection of SOM cells. Sparse sampling of this sort using single-draw Monte Carlo approaches has been shown to be quite effective at reconstructing  $N(z)$  distributions (and relevant summary statistics) to ensembles of objects, as shown in, e.g., Wittman (2009), Carrasco Kind & Brunner (2014b), and (Masters et al. 2015).

Due to the small number of occupied cells in our observed-frame SOM ( $\sim 21,600$  out of 27,000), we decide to simply draw an equal number of samples from occupied cells on the SOM. The corresponding  $P(z|\mathbf{F}_{\text{obs}})$  can then be directly calculated via

$$P(z|\mathbf{F}_{\text{obs}}) \approx \sum_{\text{cell}} \sum_{\mathcal{R}(\text{cell})} N(\mathbf{x}_{\text{som}}^{\text{cell}}) P(\mathbf{F}_{\text{model}}^{\mathcal{R}_i} | \mathbf{F}_{\text{obs}}), \quad (16)$$

where  $\sum_{\text{cell}}$  is taken over all occupied cells,  $\mathcal{R}(\text{cell})$  is a collection of  $N_{\text{MC}}$  random numbers  $\{\mathcal{R}_i\}_i$  that select the corresponding indices of model fluxes  $\mathbf{F}_{\text{model}}^{\mathcal{R}_i}$  contained within a given cell, and  $N(\mathbf{x}_{\text{som}}^{\text{cell}})$  accounts for the different occupation rates of each cell. We set  $N_{\text{MC}} = 5$  to allow for more robust sampling than single-draw approaches while still keeping the number of total draws ( $\sim 110,000$ ) a factor of  $\sim 30$  below the total number of available models.

### 3.6 SOM\_Hierarchical\_ImportanceSampling

While the stratified Monte Carlo approach is easy to implement and relatively resilient against random misses (since every cell on the SOM is sampled evenly), most objects with well-localized colors are strongly associated with a very small set of clusters. As a result, a significant amount of computation time will be wasted evaluating clusters containing model fluxes with  $P(\mathbf{F}_{\text{obs}}|\mathbf{F}_{\text{model}}) \approx 0$ . Rather than drawing an equal number of samples from every SOM cell, we could use preliminary approximations of the likelihood across the SOM to instead take advantage of adaptive *importance sampling* to favor target cells that are more likely to contain the majority of the likelihood.

In the scheme of the hierarchical interpretation outlined above in §3.5, we can approximate the likelihood contained in each individual cluster  $P(\mathbf{x}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}})$  via the identity

$$\begin{aligned} P(\mathbf{F}_{\text{obs}}) &= \int P(\mathbf{x}_{\text{som}}^{\text{cell}}) P(\mathbf{F}_{\text{obs}}|\mathbf{x}_{\text{som}}^{\text{cell}}) d\mathbf{x}_{\text{som}}^{\text{cell}} \\ &= \int \pi(\mathbf{x}_{\text{som}}^{\text{cell}}) \frac{P(\mathbf{x}_{\text{som}}^{\text{cell}}) P(\mathbf{F}_{\text{obs}}|\mathbf{x}_{\text{som}}^{\text{cell}})}{\pi(\mathbf{x}_{\text{som}}^{\text{cell}})} d\mathbf{x}_{\text{som}}^{\text{cell}}, \end{aligned} \quad (17)$$

where  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$  is a generic approximation to  $P(\mathbf{x}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}})$ . For a series of  $N_{\text{IS}}$  samples drawn from  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$ , the  $n$ th posterior moment of  $P(\mathbf{x}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}})$  can then be estimated as

$$\begin{aligned} \int (\mathbf{x}_{\text{som}}^{\text{cell}})^n P(\mathbf{x}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}}) d\mathbf{x}_{\text{som}}^{\text{cell}} &= \frac{\int (\mathbf{x}_{\text{som}}^{\text{cell}})^n \pi(\mathbf{x}_{\text{som}}^{\text{cell}}) w(\mathbf{x}_{\text{som}}^{\text{cell}}) d\mathbf{x}_{\text{som}}^{\text{cell}}}{\int \pi(\mathbf{x}_{\text{som}}^{\text{cell}}) w(\mathbf{x}_{\text{som}}^{\text{cell}}) d\mathbf{x}_{\text{som}}^{\text{cell}}} \\ &\approx \sum_{N_{\text{IS}}} (\mathbf{x}_{\text{som}}^{\text{cell}})^n w(\mathbf{x}_{\text{som}}^{\text{cell}}) / \sum_{N_{\text{IS}}} w(\mathbf{x}_{\text{som}}^{\text{cell}}), \end{aligned} \quad (18)$$

where  $w(\mathbf{x}_{\text{som}}^{\text{cell}}) = \frac{P(\mathbf{x}_{\text{som}}^{\text{cell}}) P(\mathbf{F}_{\text{obs}}|\mathbf{x}_{\text{som}}^{\text{cell}})}{\pi(\mathbf{x}_{\text{som}}^{\text{cell}})}$  is the associated *importance weight* – the overall likelihood of a given trial formed by a combination of the initial prior and the computed likelihood normalized by the proposal distribution. Estimates of the relevant marginalized distributions can be constructed by assigning each sampled location an appropriate *kernel density* with an amplitude proportional to its importance weight.

Using the same logic as §3.5, implementing a Monte Carlo sampling scheme within each cluster allows us to separate out the proposal distribution  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$  which operates on the set of clusters and the associated importance weights  $w(\mathbf{F}_{\text{model}})$  assigned to each particular chosen model. Assuming  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$  is properly normalized, the corresponding importance weight for a random model  $\mathbf{F}_{\text{model}}^{\mathcal{R}_i}$  located at cell  $\mathbf{x}_{\text{som}}^{\text{cell}}$  is then

$$w(\mathbf{F}_{\text{model}}^{\mathcal{R}_i}) = N(\mathbf{x}_{\text{som}}^{\text{cell}}) \frac{P(\mathbf{F}_{\text{model}}^{\mathcal{R}_i} | \mathbf{F}_{\text{obs}})}{\pi(\mathbf{x}_{\text{som}}^{\text{cell}})}. \quad (19)$$

For a number of  $N_{\text{IS}}$  samples drawn randomly from cells drawn from  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$ , we can then estimate  $P(z|\mathbf{F}_{\text{obs}})$  as

$$P(z|\mathbf{F}_{\text{obs}}) \approx \sum_{\mathcal{R}(\pi(\mathbf{x}_{\text{som}}^{\text{cell}}))} w(\mathbf{F}_{\text{model}}^{\mathcal{R}_i}) P(\mathbf{F}_{\text{model}}^{\mathcal{R}_i} | \mathbf{F}_{\text{obs}}). \quad (20)$$

In this paper, we have avoided the use of kernels to avoid additional computational overhead.

For the implementation used in this paper, we further build on this general approach using an adaptive importance sampling scheme equivalent to a crude form of Sequential/Population Monte Carlo (Doucet et al. 2001) and related methods:

- (i) Compute  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$  using the stratified Monte Carlo scheme outlined in §3.5 with  $N_{\text{MC}}$  random models sampled in each cell.
- (ii) Draw  $N_{\text{batch}}$  cells from  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$ .
- (iii) Draw a random model  $\mathbf{F}_{\text{model}}^{\mathcal{R}_i}$  from each cell.
- (iv) Use the corresponding set of  $w(\mathbf{F}_{\text{model}}^{\mathcal{R}_i})$ ’s to update  $\pi(\mathbf{x}_{\text{som}}^{\text{cell}})$ .
- (v) Repeat steps (ii) through (iv) for  $N_{\text{iter}}$  iterations.
- (vi) Use the total number of objects sampled from steps (i) through (v) to compute  $P(z|\mathbf{F}_{\text{obs}})$ .

We set  $N_{\text{MC}} = 1$ ,  $N_{\text{batch}} = 2500$ , and  $N_{\text{iter}} = 20$  in this paper, giving us a total of  $\sim 70,000$  models comparisons (slightly less than the amount used in §3.5). In addition, we impose a “likelihood floor” for each computed weight of  $10^{-3}$  when updating the proposal distribution both to prevent multi-valued inverse functions when drawing from the proposal and to ensure each cell has a small (but non-negligible) probability of being selected with each draw in order to ensure great robustness during sampling. Note that this likelihood floor is *not* used when constructing  $P(z|\mathbf{F}_{\text{model}})$ .

### 3.7 SOM\_CellModel\_LimitedSum

Instead of sampling over the entire SOM, we can instead attempt to limit the total number of terms included in the sum to only cells which carry significant weight, after which we can just compute the reduced set of  $P(z|\mathbf{x}_{\text{som}}^{\text{cell}})$ 's directly. In other words, after computing the full set of  $P(\mathbf{F}_{\text{obs}}, \mathbf{F}_{\text{som}}^{\text{cell}})$ 's across the 3-D observed-frame SOM for a particular object, we select a collection of  $k$  cells such that

$$P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}}) > c_{\eta} \max \left[ P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}}) \right], \quad (21)$$

where  $c_{\eta}$  is a specified acceptance threshold which we take to be  $10^{-3}$  and  $\mathbf{F}_{\text{som}}^{\text{cell}}$  is the associated SOM cell model (see §2.2), which we assume to be a reasonable proxy of the data in a given cell. As noted in §3.5, while this assumption is not always true (especially in sparsely populated and/or highly variable regions of color space), it is a good first-order approximation to the problem.

After using the cell models to select our best-fitting cells, we can then compute  $P(z|\mathbf{F}_{\text{obs}})$  as

$$P(z|\mathbf{F}_{\text{obs}}) \approx \sum_k P(z|\mathbf{x}_k) = \sum_k \sum_{\text{dust}}^{\mathbf{x}_k} \sum_{\text{gal}}^{\mathbf{x}_k} \sum_{\phi}^{\mathbf{x}_k} P(\theta, \phi|\mathbf{F}_{\text{obs}}), \quad (22)$$

where  $\sum^{\mathbf{x}_k}$  indicates the sum is taken over all objects sorted into the  $k$ th cell. As the number of terms included in the sum over  $k$  should only be a small fraction of the total number of cells, it should be reasonably quick to compute.

### 3.8 SOM\_CellModel\_Average

Instead of comparing individual models and/or fuzzy templates, we now wish to take advantage of the quantization the SOM provides by assigning the models to a discrete number of cells. In particular, every cell in our 3-D observed-frame SOM has an associated redshift number distribution  $N(z|\mathbf{x}_{\text{som}}^{\text{cell}})$ , whose combined sum would be proportional to the input redshift prior which all the objects in our model catalog were drawn from. For any individual object, this  $N(z|\mathbf{x}_{\text{som}}^{\text{cell}})$  distribution corresponds to a specific redshift probability  $P(z|\mathbf{x}_{\text{som}}^{\text{cell}}, \mathbf{F}_{\text{obs}})$  for the given cell through a redshift-dependent weight function  $w(z|\mathbf{x}_{\text{som}}^{\text{cell}}, \mathbf{F}_{\text{obs}})$ . This allows us to write  $P(z|\mathbf{F}_{\text{obs}})$  as

$$P(z|\mathbf{F}_{\text{obs}}) = \sum_{\text{cell}} P(z|\mathbf{x}_{\text{som}}^{\text{cell}}, \mathbf{F}_{\text{obs}}) = \sum_{\text{cell}} w(z|\mathbf{F}_{\text{obs}}, \mathbf{x}_{\text{som}}^{\text{cell}}) N(z|\mathbf{x}_{\text{som}}^{\text{cell}}), \quad (23)$$

where unoccupied cells again are excluded from the sum.

Assuming that the SOM has clustered objects efficiently relative to the observed errors on  $\mathbf{F}_{\text{obs}}$ , we roughly can approximate the likelihood of a given cell as simply being proportional to its intrinsic redshift distribution such that  $P(z|\mathbf{F}_{\text{obs}}, \mathbf{x}_{\text{som}}^{\text{cell}}) \sim N(z|\mathbf{x}_{\text{som}}^{\text{cell}})$ . If we further assume (as in §3.7) that the cell model is a good proxy for the data in a given cell, we can further approximate the cell weight as being the likelihood of the corresponding cell model  $P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}})$ . This gives, up to a normalization factor,

$$P(z|\mathbf{F}_{\text{obs}}) \approx \sum_{\text{cell}} P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}}) N(z|\mathbf{x}_{\text{som}}^{\text{cell}}). \quad (24)$$

As each  $N(z|\mathbf{x}_{\text{som}}^{\text{cell}})$  is an array of  $N_z$  elements, this sum

is computationally taxing due to the large number of cells used in the construction of the observed-frame SOM. As a result, we instead opt to use a point estimate for  $N(z|\mathbf{x}_{\text{som}}^{\text{cell}})$  based on its zeroth moment (i.e. mean)  $\bar{z}(\mathbf{x}_{\text{som}}^{\text{cell}})$ , which can be pre-computed and saved to memory. The mean weighted redshift across the SOM for a given object can then be written as

$$z_{\text{avg}}(\mathbf{F}_{\text{obs}}) = \frac{\sum_{\text{cell}} P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}}) \bar{z}(\mathbf{x}_{\text{som}}^{\text{cell}})}{\sum_{\text{cell}} P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}})}. \quad (25)$$

While the full  $P(z|\mathbf{F}_{\text{model}})$  distribution could be recovered using kernel density estimation centered on the corresponding collection of  $\bar{z}(\mathbf{x}_{\text{som}}^{\text{cell}})$ 's and  $P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}})$ 's (see §3.6), for computational convenience we instead opt to compute only the first moment

$$\sigma_{z,\text{avg}}^2(\mathbf{F}_{\text{obs}}|z_{\text{avg}}) = \frac{\sum_{\text{cell}} P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}}) [\bar{z}(\mathbf{x}_{\text{som}}^{\text{cell}}) - z_{\text{avg}}(\mathbf{F}_{\text{obs}})]^2}{\sum_{\text{cell}} P(\mathbf{F}_{\text{som}}^{\text{cell}}|\mathbf{F}_{\text{obs}})}. \quad (26)$$

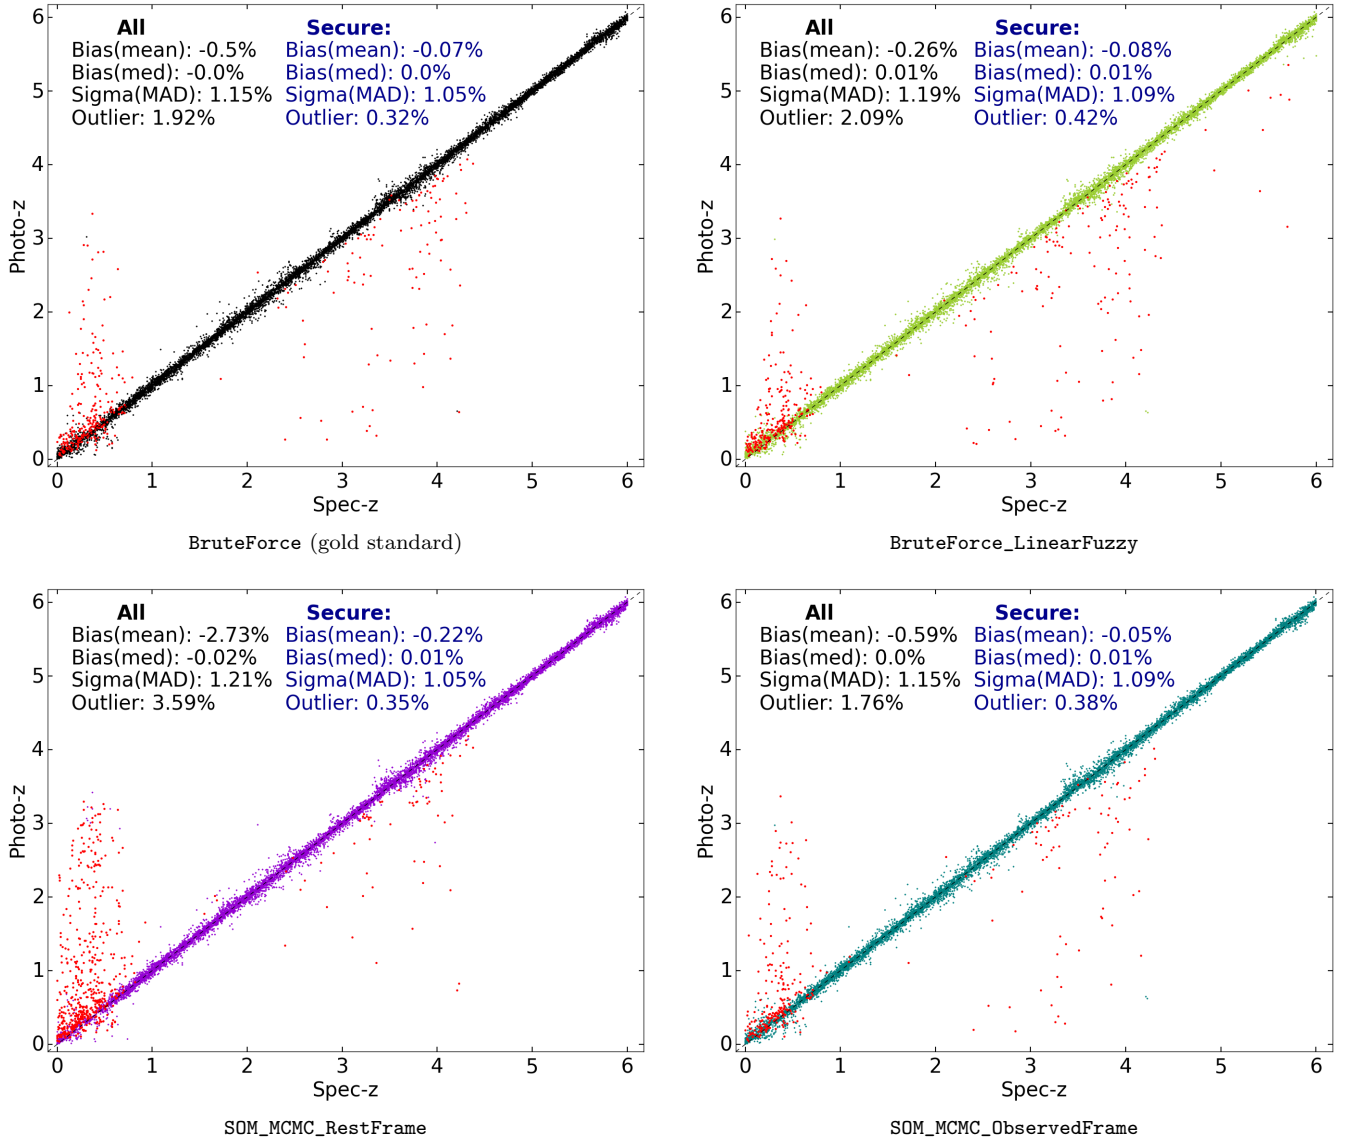
Given that this method is the “crudest” approach that makes the strongest assumptions concerning the data and the SOM, we take this as a suitable estimate of the shape of the corresponding PDF.

### 3.9 Summary

We have outlined eight separate methods to derive photo-*z*'s for individual objects. Each pair of methods (brute-force, MCMC, hierarchical, cell model-dependent) represent four different ways to use templates to derive  $P(z|\mathbf{F}_{\text{obs}})$ . In our brute-force approaches, all available models and/or model parameters are compared to the data in order to derive the appropriate redshift PDF. In our MCMC-based approaches, we attempt to explore the set of models contained within contiguous regions on the SOM. In our hierarchical sampling approaches, we randomly sample galaxies according to the SOM's partitions in color space to stochastically construct estimates to the specific portion of the SOM-based color-redshift relation for each individual galaxy. Finally, in our cell model-dependent approaches, we attempt to “triage” the space spanned by the SOM using a series of summary statistics in order to derive limited estimates of  $P(z|\mathbf{F}_{\text{obs}})$ . Each take advantage of different aspects of our initial template set and the final SOM outputs and represent a variety of ways to take advantage of the SOM in order to maximize the utility of our fuzzy templates.

## 4 PERFORMANCE TESTS AGAINST MOCK LSST AND EUCLID DATA

Although we would ideally like to test each of these approaches using real data, preliminary tests with a number of individual objects taken from Cosmological Origins Survey (COSMOS; Scoville et al. 2007) data indicate that the limited range spanned by portions of the Brown et al. (2014) archetypes are unrepresentative of a number of actively star-forming galaxies at higher redshift. As noted in Paper I, similar issues exist among all currently used template sets



**Figure 1.** The input ( $z_{\text{spec}}$ ) versus fitted ( $\bar{z}_{\text{phot}}$ ) redshifts for our BruteForce (our “gold standard”; black), BruteForce\_LinearFuzzy (yellow-green), SOM\_MCMC\_RestFrame (purple), and SOM\_MCMC\_ObservedFrame (teal) methods (see §3) to the 10,000 objects in our 9-band LSST *ugrizY* and *Euclid YJH* mock catalog consisting of  $Y_{\text{LSST}} = 24$  mag objects.  $P(z|\mathbf{F}_{\text{obs}})$  distributions with redshift-normalized standard deviations  $\sigma'_z > 0.15$  are flagged as insecure and plotted as red points. The redshift-normalized mean bias  $\bar{\Delta z}'$ , median bias  $\Delta z'_{50}$ ,  $1\sigma$  median absolute deviation (MAD)  $\sigma'_{z,\text{MAD}}$ , and the catastrophic outlier fraction  $f_{\text{cat}}$  for the full and secure samples are listed in each plot. See Figure 2 for more details.

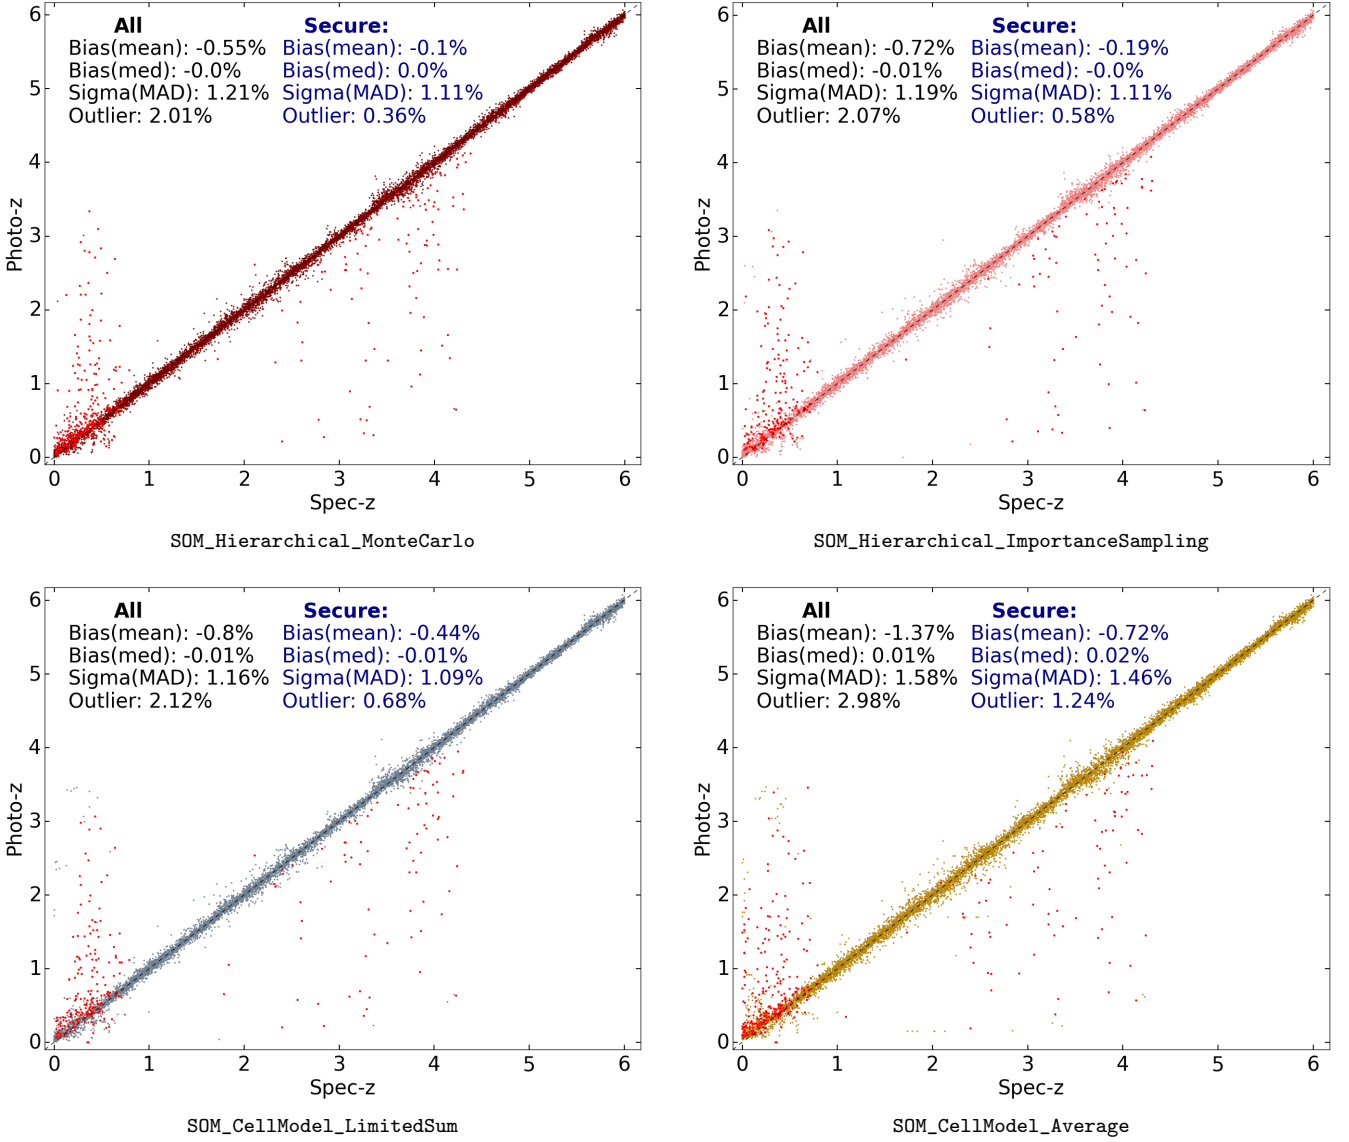
(Coleman et al. 1980; Kinney et al. 1996; Polletta et al. 2007). Rather than supplementing the Brown et al. (2014) galaxies with synthetic spectra created from stellar population synthesis models, we instead use a mock catalog drawn from our fuzzy templates to serve as performance benchmarks for our approach. While this does not give us a fully representative idea of how each of our individual approaches work in practice, it does give us an idealized comparison as to their *intrinsic effectiveness* assuming no template mismatch and similar modeling assumptions, which can help to identify the most useful approaches out of the eight proposed here.

We create a mock catalog of 10,000 objects scaled to 24 mag in the LSST Y-band to examine performance at a

reasonable given magnitude. We generate each object by sampling uniformly from  $\theta$  over the redshift interval  $z = 0 - 6$  ( $\Delta z = 0.01$ ) and the given set of priors  $P(\phi)$  on our fiducial set of templates (see §2.1). Each object is then assigned Gaussian errors (added in quadrature) according to the estimated calibration uncertainty (1%) and observational errors based on the expected imaging depths, and the photometry is jittered according to the error bars.

We use each of the methods detailed in §3 to derive photo- $z$ 's to the entire mock catalog. Although we save the full  $P(z|\mathbf{F}_{\text{obs}})$  for each object when available (i.e. excluding SOM\_CellModel\_Average), we choose to use the zeroth moment  $\bar{z}_{\text{phot}}$  and the square root of the first moment  $\sigma_z$  as our fiducial photo- $z$  point estimates and their associated errors.





**Figure 2.** As Figure 1, but for our SOM\_Hierarchical\_MonteCarlo (dark red), SOM\_Hierarchical\_ImportanceSampling (pink), SOM\_CellModel\_LimitedSum (gray), and SOM\_CellModel\_Average (dark yellow) methods (see §3). While all of our methods are able to reproduce the one-to-one relationship between spec-z and photo-z quite well with  $\sigma'_{z,\text{MAD}} \lesssim 0.015$ , our SOM\_CellModel\_LimitedSum and SOM\_CellModel\_Average approaches consistently fail to identify unreliable photo-z's due to subtle biases introduced from their reliance on SOM cell models. Once insecure photo-z's have been removed, we find the majority of methods achieve accuracies comparable to our brute-force approaches with  $|\Delta z'| \approx 0.5\%$ ,  $\sigma'_{z,\text{MAD}} \approx 1.1\%$ , and catastrophic outlier rates of  $\lesssim 0.4\%$ , with our BruteForce\_LinearFuzzy, SOM\_MCMC\_ObservedFrame, and SOM\_Hierarchical\_MonteCarlo giving the best overall results.

We report five summary statistics for each sample: the redshift-normalized *mean bias*,

$$\Delta z' = \left\langle \frac{|z_{\text{spec}} - \bar{z}_{\text{phot}}|}{1 + z_{\text{spec}}} \right\rangle, \quad (27)$$

the redshift-normalized *median bias*,

$$\Delta z'_{50} = \text{median} \left( \frac{|z_{\text{spec}} - \bar{z}_{\text{phot}}|}{1 + z_{\text{spec}}} \right), \quad (28)$$

the redshift-normalized *mean scatter*,

$$\sigma(\Delta z') = \left[ \text{Var} \left( \frac{|z_{\text{spec}} - \bar{z}_{\text{phot}}|}{1 + z_{\text{spec}}} \right) \right]^{-\frac{1}{2}} \quad (29)$$

the  $1\sigma$  *median absolute deviation* (MAD),


$$\sigma'_{z,\text{MAD}} = 68.3 \text{ percentile} \left( \left| \frac{z_{\text{spec}} - \bar{z}_{\text{phot}}}{1 + z_{\text{spec}}} - \Delta z'_{50} \right| \right), \quad (30)$$

and the *catastrophic outlier fraction*,

$$f_{\text{cat}} = N(\Delta z' > 0.15) / N. \quad (31)$$

$z_{\text{spec}}$  in all cases is taken to be the input redshift, while  $\bar{z}_{\text{phot}}$  is the estimated redshift from each of the different methods outlined in §3.

In order to remove multimodal and/or unconstrained fits, we classify our objects into “secure” and “insecure” redshifts based on their mean redshift-normalized *standard de-*



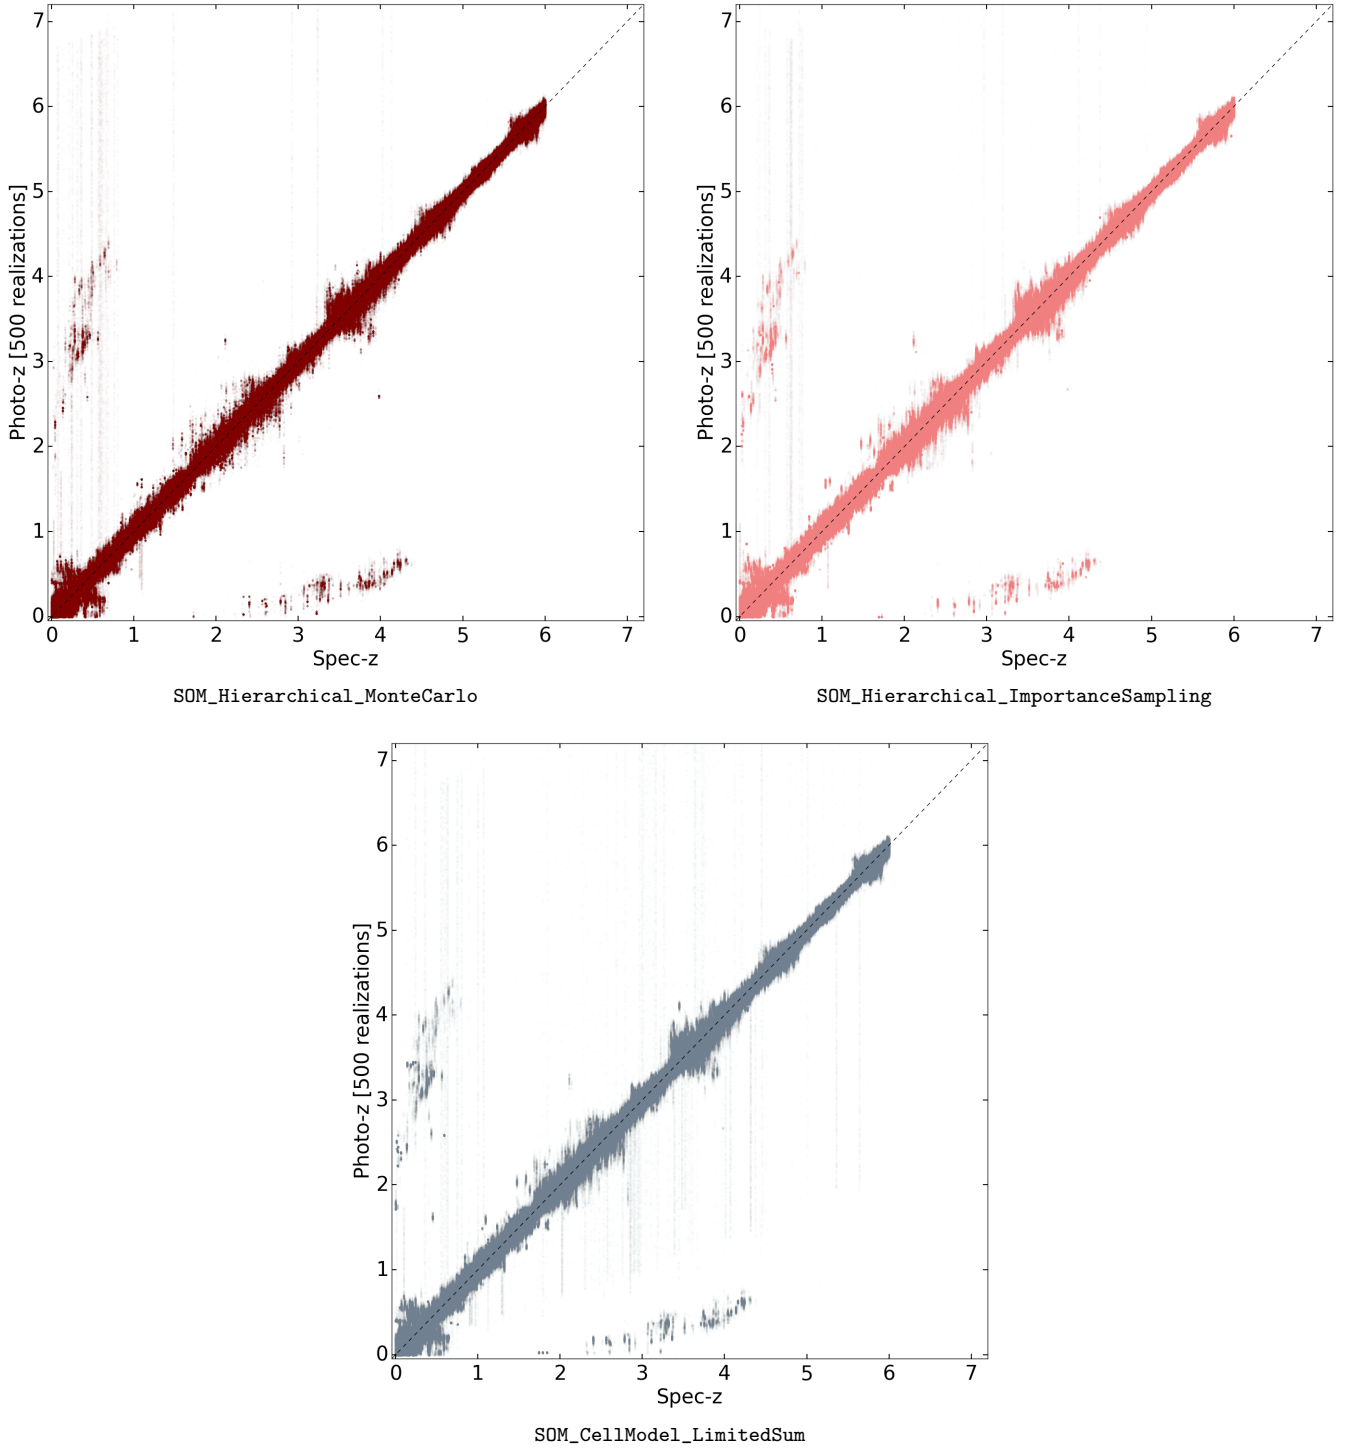
**Figure 3.** The two-dimensional stacked  $P(z|\mathbf{F}_{\text{obs}})$  distributions for our BruteForce (gold standard), BruteForce\_LinearFuzzy, SOM\_MCMC\_RestFrame, and SOM\_MCMC\_ObservedFrame methods (see §3) generated using 500 Monte Carlo samples drawn from the redshift PDF of each object and plotted using the same color scheme as Figure 1. See Figure 4 for more details.

violation,

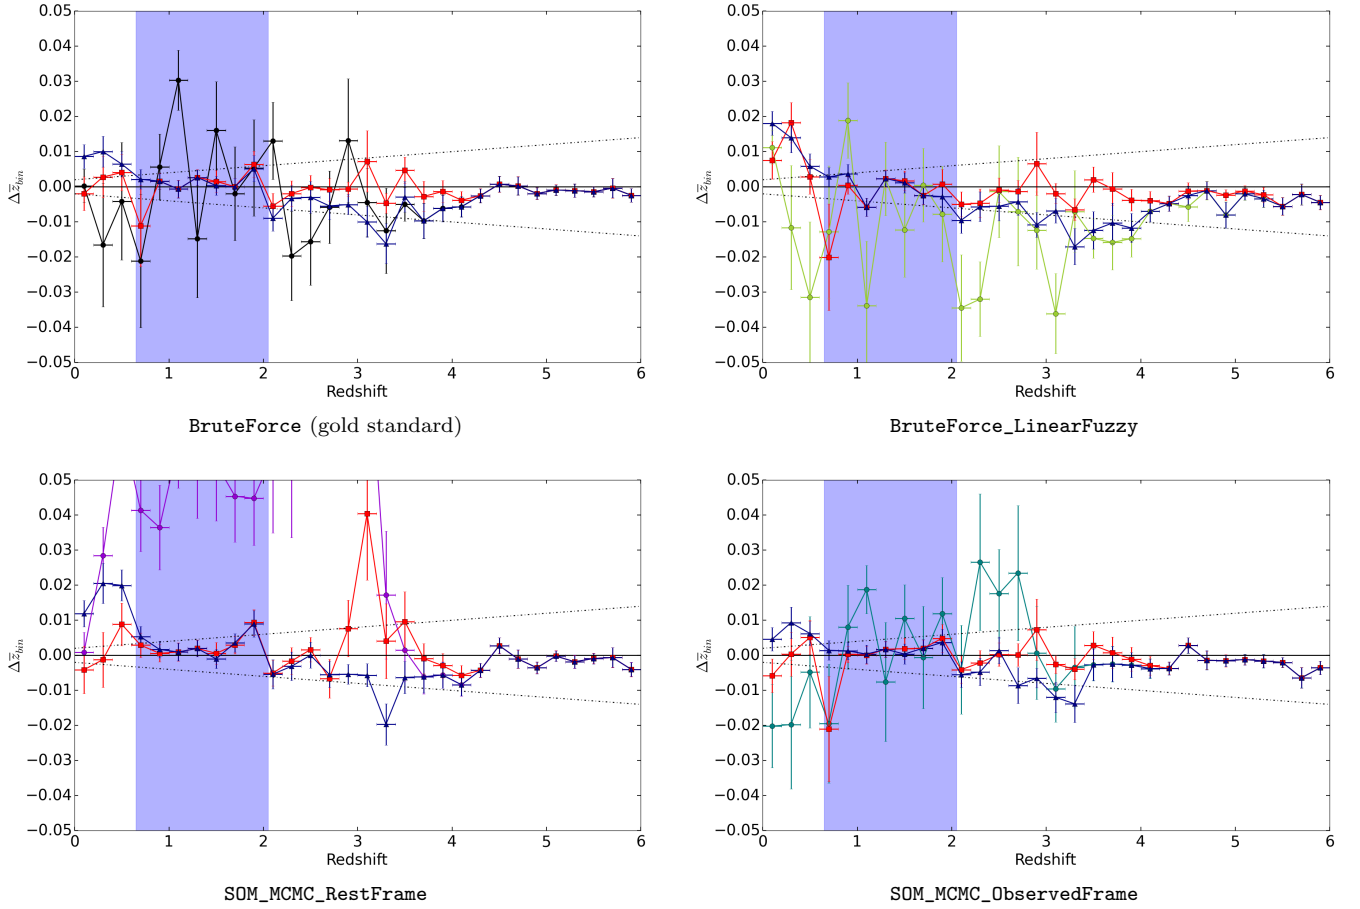
$$\sigma'_z = \sigma_z / (1 + \bar{z}). \quad (32)$$

Objects are flagged as insecure if  $\sigma'_z > 0.15$  (i.e. if the standard deviation of the redshift PDF is larger than the catastrophic error boundaries), and otherwise are classified as se-

cure. These are then used to construct “full” and “secure” samples. This is done to ensure that we have an unbiased view of the effectiveness of our redshift methods to not only derive accurate photo- $z$ ’s but also to flag outliers. As weak lensing approaches are extremely sensitive to such outliers (Cunha et al. 2014), being able to accurately detect and re-



**Figure 4.** As Figure 3, but for our `SOM_Hierarchical_MonteCarlo`, `SOM_Hierarchical_ImportanceSampling`, and `SOM_CellModel_LimitedSum` methods (see §3). All methods recover the general degeneracies present across the objects in our catalog – including redshift-reddening degeneracies and confusion between the 1216 Å and 4000 Å breaks. We find that the increased flexibility of `BruteForce_LinearFuzzy` not only sharpens degeneracies as compared to `BruteForce` but also creates new ones, such as a thin band of probability at lower redshift for sources at  $z \sim 4-6$ . In addition, we confirm that the large band of flagged outliers at low redshift from `SOM_MCMC_RestFrame` (see Figure 1) is due to an overabundance of probability at higher redshifts for low- $z$  sources, illustrating the impact of the redshift gradient noted in Speagle & Eisenstein (2015a). Finally, while `SOM_Hierarchical_MonteCarlo` gives good overall performance, its PDFs are noticeably “rougher” than the other approaches, likely due to sparse sampling in the small fraction of cells where the models with the highest likelihoods are contained.



**Figure 5.** The mean redshift bias  $\Delta\bar{z}_{\text{bin}} = \langle z_{\text{spec}} \rangle_{\text{bin}} - \langle \bar{z}_{\text{phot}} \rangle_{\text{bin}}$  in a given  $\bar{z}_{\text{phot}}$ -selected redshift tomographic bin, plotted as a function of redshift for our **BruteForce** (gold standard), **BruteForce\_LinearFuzzy**, **SOM\_MCMC\_RestFrame**, and **SOM\_MCMC\_ObservedFrame** methods (see §3). The results for the full sample are colored using the same color scheme as Figure 1, while those from the secure subsample are shown in red. For comparison, the results from a “cleaned” subsample where all catastrophic outliers (i.e. objects with  $\Delta z' > 0.15$ ) have been removed are shown in blue. Errors for each sample have been derived through bootstrap resampling, with  $1\sigma$  errors plotted. The *Euclid* photo- $z$  accuracy requirements for weak-lensing of  $\Delta\bar{z} < 0.002(1+z)$  from [Laureijs et al. \(2011\)](#) are over-plotted as dotted-dashed black lines, with the relevant redshift range highlighted in light blue. See Figure 6 for more details.

move badly constrained photo- $z$ ’s are key to their successful application.<sup>4</sup>

The generally quality of the photo- $z$  point estimates derived from each method is illustrated in Figures 1 and 2, along with several summary statistics. We find that all of the methods are able to recover the general one-to-one relation between the input and output redshift relatively well with  $\sigma'_{z,\text{MAD}} \lesssim 1.15\%$ . However, we find that both our **SOM\_CellModel\_LimitedSum** and **SOM\_CellModel\_Average** methods are unable to consistently flag insecure photo- $z$ ’s, impeding their ability to remove catastrophic outliers that can severely bias weak lensing analyses. We find this is mainly due to the assumption that the SOM cell models are a reliable proxy for the data clustered around them:

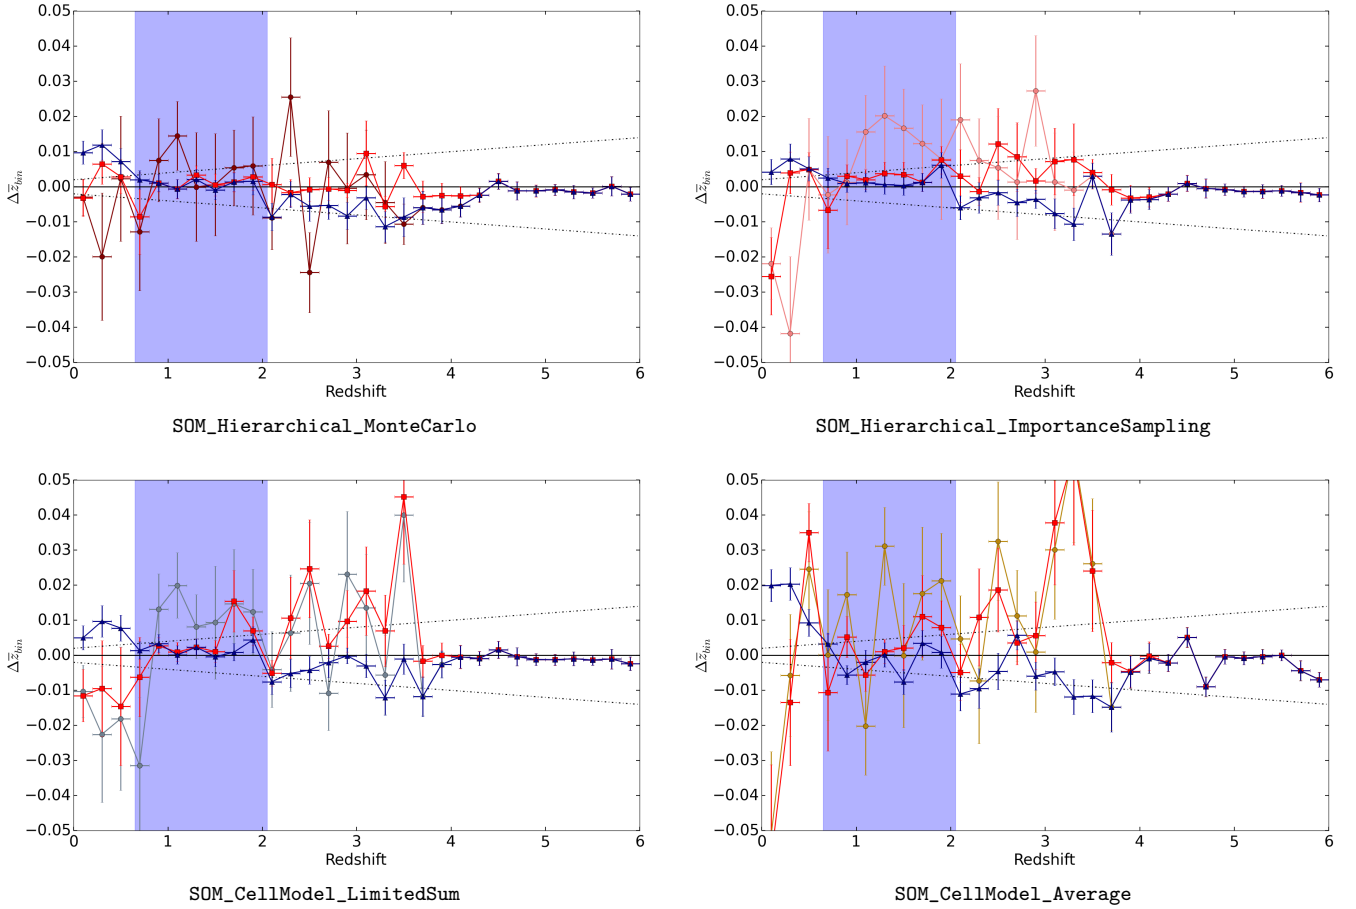
<sup>4</sup> We note that median-based point estimates  $z_{50,\text{phot}}$  give comparable results to  $\bar{z}_{\text{phot}}$  in terms of overall accuracy. However, median-based estimates of the relative spread  $\sigma_{z,\text{MAD}}$  are less effective at flagging possible catastrophic errors because the majority of the multi-modal PDFs tends to be concentrated near one mode.

while this assumption is true in general, it introduces a color-dependent bias during the sampling process that distorts the derived redshift PDF. Even in the worst case, however, we find that our catastrophic outlier fraction is  $\lesssim 3.5\%$ .

After limiting each of our other six methods to their secure subsamples, we find their accuracy noticeably improves, with redshift-normalized mean biases  $\lesssim 0.1\%$  and catastrophic outlier fractions of  $\sim 0.35\%$ , a reduction of a factor of  $\sim 5$ – $10$ . While both **SOM\_CellModel\_LimitedSum** and **SOM\_CellModel\_Average** also display improved performance, they are only able to reduce their catastrophic outlier fractions by a factor of 2–3.

Among our ensemble of methods, we find that **BruteForce\_LinearFuzzy**, **SOM\_MCMC\_ObservedFrame**, and **SOM\_Hierarchical\_MonteCarlo** give the best results relative to **BruteForce** before and after the removal of insecure redshifts. In addition, although our **SOM\_Hierarchical\_ImportanceSampling** approach gives a slightly higher catastrophic outlier fraction after removing photo- $z$ ’s flagged as insecure, we find its general overall performance is also similar to **BruteForce**. Finally, although






**Figure 6.** As Figure 5, but for our `SOM_Hierarchical_MonteCarlo`, `SOM_Hierarchical_ImportanceSampling`, `SOM_CellModel_LimitedSum`, and `SOM_CellModel_Average` methods (see §3). For the full sample, numerous outliers lead to significant biases for most redshift bins below  $z \lesssim 4$ . While variation of  $\Delta\bar{z}_{\text{bin}}(z_{\text{bin}})$  appears to be symmetrical for the `BruteForce` approach and the sparse-sampling `SOM_Hierarchical_MonteCarlo` approaches, in the majority of cases the overabundance of up-scattered redshifts tend to drive variation in  $\Delta\bar{z}_{\text{bin}}(z_{\text{bin}})$  to be positive (with the exception of `BruteForce_LinearFuzzy`, which introduces new low- $z$  degeneracies; see Figure 3). After restricting to secure photo- $z$ 's only, the bias decreases substantially (except for `SOM_CellModel`-based approaches; see Figure 2), with the majority of tomographic bins for the secure subsample meeting the *Euclid* photo- $z$  accuracy requirements above  $z \sim 0.8$ .

our `SOM_MCMC_RestFrame` run displays a noticeably high concentration of misfit objects scattered from low to high redshift, when comparing only the secure subsamples we find its performance is nearly identical to the previous two methods. We thus find both our MCMC-driven and hierarchical sampling methods can give comparable results relative to a brute-force search and are effective at locating and characterizing  $P(z|\mathbf{F}_{\text{obs}})$ .

To examine the redshift PDFs from our individual methods more closely, in Figure 3 we plot their two-dimensional stacked  $P(z|\mathbf{F}_{\text{obs}})$  distributions using 500 Monte Carlo realizations drawn from the redshift PDF of each object. As can be clearly seen, all of our approaches recover the broad features and degeneracies observed in the `BruteForce` run, including the low- $z$  redshift-reddening degeneracy (as well as several others at higher redshifts) as well as confusion between the 1216 Å and 4000 Å breaks. We confirm that the `SOM_MCMC_RestFrame` case displays an overabundance of objects and corresponding probability at higher redshifts, which is likely due to the impact of the

redshift gradient observed in first noted in Paper I, which tends to increase the likelihood that chains converge to higher redshift solutions. In addition, we find that while our `BruteForce_LinearFuzzy` run performs comparably to `BruteForce` run, the increased flexibility actually causes existing degeneracies to sharpen while simultaneously creating new degeneracies not present in any of the other runs. This indicates that while marginalizing over the linear parameters of our fuzzy templates is generally effective, it occasionally can enhance the likelihoods of degenerate solutions.

In order to measure the dark energy equation-of-state, LSST and *Euclid* will rely on weak lensing to determine the cosmic shear in given redshift tomographic bins to sub-percent accuracy (Ivezic et al. 2008; Laureijs et al. 2011). This approach depends on having extremely accurate measurements of the mean redshift of a given bin. To illustrate the accuracy of each approach more clearly in the context of measuring cosmic shear, we plot the evolution in mean bias  $\Delta\bar{z} = \langle z_{\text{spec}} \rangle_{\text{bin}} - \langle z_{\text{phot}} \rangle_{\text{bin}}$  and the normalized standard deviation  $\sigma'_z$  in fixed  $\Delta z = 0.2$  redshift tomographic



**Figure 7.** As Figure 5, but for the redshift-normalized mean scatter  $\sigma(\Delta z')$  for our `BruteForce`, `BruteForce_LinearFuzzy`, `SOM_MCMC_RestFrame`, and `SOM_MCMC_ObservedFrame` methods (see §3). The *Euclid* photo- $z$  accuracy requirements for weak-lensing of  $\sigma'_z < 0.05$  from [Laureijs et al. \(2011\)](#) are over-plotted as dotted-dashed black lines, with the relevant redshift range highlighted in light blue. See Figure 8 for more details.

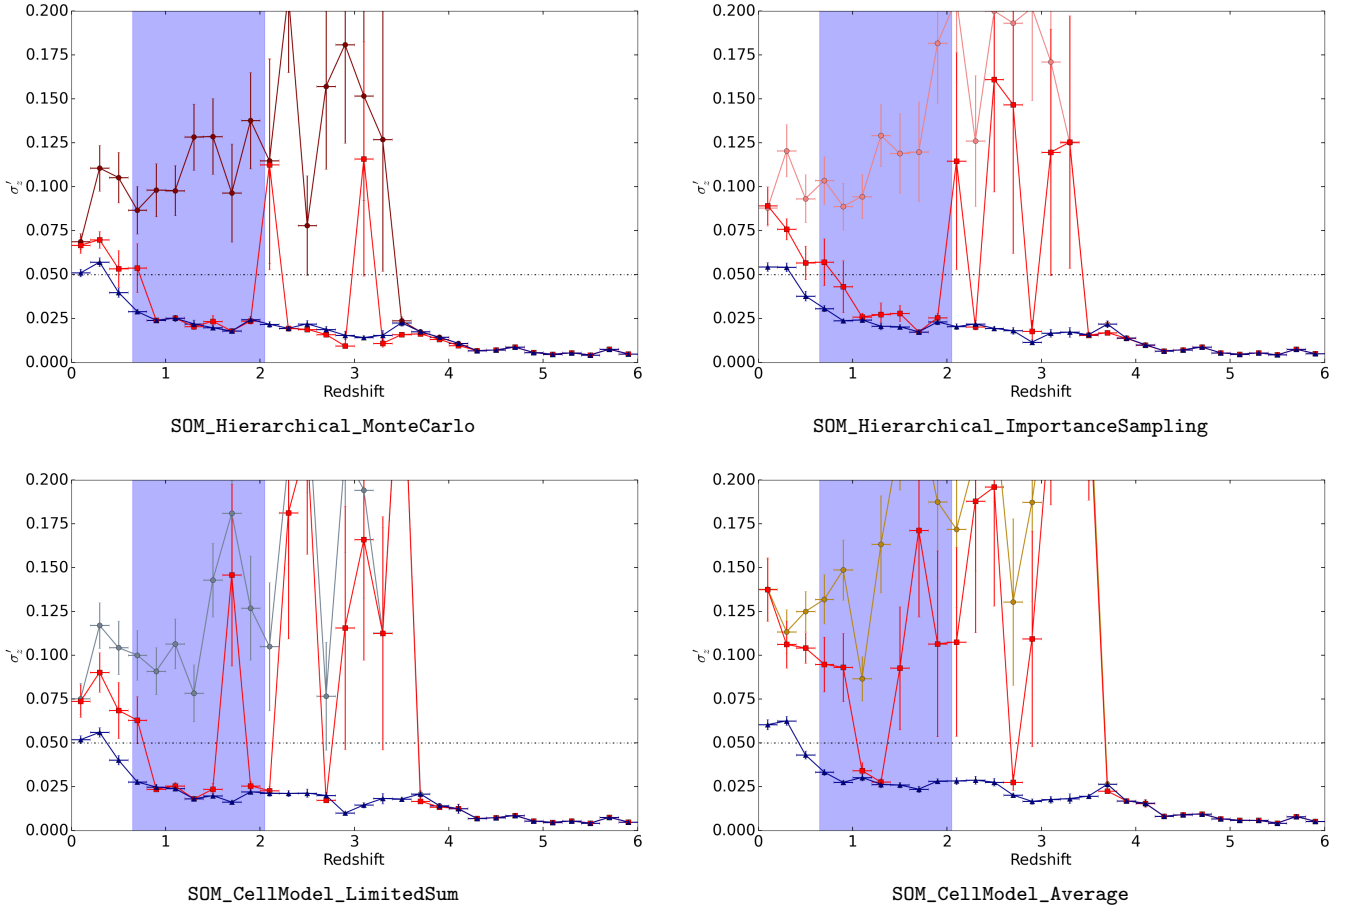
bins (selected using  $z_{\text{phot}}$ ) from  $z = 0-6$  in Figures 5, 6, 7, and 8. To establish an appropriate baseline comparison for each method, we compare both the full and secure samples to an explicitly “cleaned” sample where all catastrophic outliers (i.e. objects with  $\Delta z' > 0.15$ ) have been removed.

We find that for all of our methods (excluding our `SOM_CellModel`-based approaches), the secure subsamples are able to meet the *Euclid* mean photo- $z$  bias accuracy requirements for weak-lensing of  $\Delta \bar{z} < 0.002(1+z)$  for most redshift bins above  $z \sim 0.8$ . Without removing insecure objects, however, we find the full sample fails to meet the accuracy requirements for almost all redshift bins where  $z \lesssim 4$ . After this point, the redshift identification for most objects becomes more secure due to the prominence of the 1216 Å break combined with *Euclid*’s NIR coverage in the *YJH* bands.

Likewise, the observed redshift-normalized scatter  $\sigma'_z$  for the secure subsample successfully meets the *Euclid* photo- $z$  accuracy requirements of  $\sigma'_z < 0.05$  for most redshift bins above  $z \sim 1$  with a comfortable margin for error (although between  $z \sim 3-4$  there is some tension seen in our `SOM_MCMC_RestFrame` and `SOM_Hierarchical_ImportanceSampling` results), while the full sample fails for  $z \lesssim 3.6$ , regardless of the method used.

Together, these results illustrate the power of using the clustering methods such as the SOM to enhance photo- $z$  searches, as well as the danger in relying heavily on summary statistics rather than the underlying collection of model photometry. Ultimately, our findings suggest that the intrinsic ability of the SOM to capture variable information content and widely-separated/broad degeneracies in parameter space and project it into a reduced-dimensional manifold in a topologically smooth fashion lends a significant amount of flexibility to the model-fitting process. This feature is clearly apparent in our 3-D `SOM_MCMC_ObservedFrame` implementation, which displays performance noticeably better performance than the 2+1-D `SOM_RestFrame` counterpart that is on par with our `BruteForce` “gold standard”. In addition, our hierarchical sampling approaches demonstrate the more generally ability granted by clustering methods to partition the large amount of model photometry into useful partitions in color space that can be quickly explored.

However, while our results emphasize the SOM is a useful tool for clustering and re-organizing the space, using simple summary statistics that describe individual SOM cells to derive photo- $z$  estimates is inherently risky. As individual cell models (and cell means) can display significant variation from the model data assigned to them, either due to high



**Figure 8.** As Figure 7, but for our `SOM_Hierarchical_MonteCarlo`, `SOM_Hierarchical_ImportanceSampling`, `SOM_CellModel_LimitedSum`, and `SOM_CellModel_Average` methods (see §3). Similar to Figure 5, numerous outliers present in the full sample lead to significant biases for most redshift bins below  $z \lesssim 3.6$ , while the majority of tomographic bins for our secure photo- $z$ 's above  $z \sim 0.8$  meet the accuracy requirements (with the exception of the `SOM_CellModel`-based approaches). Due to our small sample size and the sensitivity of  $\sigma(\Delta z')$  to outliers, isolated failures in tomographic bins at  $z > 0.8$  are often due to only a small number of outliers.

intrinsic variability, sparse sampling, or both (Paper I), relying on them directly will almost guarantee that certain regions of color space will become distorted during the sampling process. This can lead to occasional biases in the final photo- $z$  estimates that make it more difficult to estimate the quality of the final results. While these approaches are reasonable for deriving quick photo- $z$  estimates, they should not be used in place of more rigorous approaches if identifying and removing catastrophic errors from the overall sample are important. It thus remains important to use these approaches judiciously, especially when abstracting away from the underlying set of models.

In addition, it is important to note that all the results derived here have utilized the same set of templates and modeling assumptions, and thus have avoided possible systematic errors due to template mismatches with observed photometry. This will most likely increase the mean bias, scatter, and catastrophic outlier fraction within a given tomographic redshift bin in a redshift-dependent manner. It is thus important to recognize these results as emulating the best-case performance of each of these sampling techniques as applied to upcoming large-scale surveys.

## 5 CONCLUSION

Due to the extremely difficult task of obtaining spectroscopic redshifts (spec- $z$ 's) to the billions of objects that will be observed in future large-scale photometric surveys, fast and accurate photometric redshifts (photo- $z$ 's) remain a necessary tool to enable “big data” science. Although significant progress has been made using machine learning to derive photo- $z$ 's (Carrasco Kind & Brunner 2013, 2014a; Sadeh et al. 2015; Bonnett 2015; Almosallam et al. 2015; Hoyle 2015), improvements in template-fitting methods have in general lagged behind, with codes today still suffering from many of the same model deficiencies and computational issues that plagued them years ago. While template-fitting methods remain effective (see Hildebrandt et al. 2010; Dahlen et al. 2013; Sánchez et al. 2014 for recent comparisons and Cool et al. 2013; Ilbert et al. 2013; Tomczak et al. 2014 for recent applications), they are in general under-equipped to handle the enormous number of objects LSST, *Euclid*, and other similar surveys are expected to observe. Because template-fitting methods will likely continue to serve a necessary function in extragalactic science for the upcoming decade, this issue is worrisome: while ma-

chine learning approaches are effective within the bounds of their training sets (Sánchez et al. 2014), the limited training set provided by spec-*z*'s today (Masters et al. 2015) and the difficulty of getting reliable spectra to higher redshift sources essentially guarantees that template-based photo-*z*'s will continue to serve a crucial role in high-*z* science.

We have attempted to test the performance of the new hybrid template-fitting and machine-learning framework outlined in Paper I that alleviates both of these issues by allowing a large set of “fuzzy archetypes” to be fit *en masse* in approximately constant time. While the observational requirements are more intensive than those used by current techniques, our method significantly reduces existing model dependencies and takes advantage of machine learning techniques such as Self Organizing Maps (SOMs) to improve the scope of parameter space that can be quickly yet robustly explored by different sampling techniques.

Using a mock set of galaxies with LSST *ugrizY* and *Euclid* *YJH* photometry with associated calibration uncertainties and observational errors, we quantify the performance of eight different approaches to deriving photo-*z*'s – a standard brute-force approach (**BruteForce**), a fuzzy template-based brute-force approach (**BruteForce\_LinearFuzzy**), two approaches based on Markov Chain Monte Carlo (MCMC) sampling over the SOM (**SOM\_MCMC\_ObservedFrame** and **SOM\_MCMC\_RestFrame**), two approaches based on hierarchical sampling over the SOM (**SOM\_Hierarchical\_MonteCarlo** and **SOM\_Hierarchical\_ImportanceSampling**), and two “quick estimate” approaches based on summary statistics derived from the SOM (**SOM\_CellModel\_LimitedSum** and **SOM\_CellModel\_Average**). We find that all of our approaches give good photo-*z* estimates with a redshift-normalized  $1\sigma$  median absolute deviation (MAD) of  $\sigma_{z,\text{MAD}} \approx 1.1\%$  and small catastrophic outlier fractions of  $\sim 2\%$ . Although we find our **SOM\_CellModel**-based approaches are not able to consistently flag insecure photo-*z*'s, our remaining MCMC-driven methods are able to robustly flag outliers and recover the underlying PDF relative to our brute-force approaches, with the results of **SOM\_MCMC\_ObservedFrame** and **SOM\_Hierarchical\_MonteCarlo** virtually indistinguishable from our **BruteForce** “gold standard”.

In addition, we find our methods are able to meet the stringent *Euclid* photo-*z* accuracy requirements for the weak lensing analysis of cosmic shear in the majority of redshift bins in our sample over  $z = 0-6$ , both in terms of the mean redshift bias as well as the general accuracy over the corresponding ensemble of objects. We thus confirm that the fundamental framework is applicable to future photo-*z* applications to wide-field surveys and can generate accurate predictions assuming a representative sample of underlying fuzzy archetypes can be obtained/generated.

There are a number of future applications of this work, some of which have been suggested in Paper I and others which have been partially explored by previous authors. In particular, while our **SOM\_CellModel**-based approaches can be computed extremely rapidly, their dependence on individual SOM cell models tends to bias the results for a number of objects. Some of this dependence might be substantially reduced or even removed with a recourse to a “Random Atlas” of SOMs (Carrasco Kind & Brunner 2014a), which we could use to derive a combined ensemble prediction as outlined in Carrasco Kind & Brunner (2014a). Alternately, we

could move away from discrete mappings of objects onto their best-matching cells onto the SOM (§2.2) in favor of general PDFs over the entire collection of SOM cells. These could (along with some set of summary statistics) then be used as a type of non-linear dimensionality reduction whose components would be fed into other machine-learning architectures in an attempt to predict the color-redshift relationship at a given location.

Finally, one could investigate taking the general framework outlined and tested here and applying it to the large regions of parameter space spanned by stellar population synthesis models in an effort to achieve similar results. As the large numbers of parameters used in these models render direct Monte Carlo sampling methods intractable, such an approach would have to involve iterative training steps where a small amount of objects (fit using more computationally expensive methods) would be used to approximate the manifold (and degeneracies) where galaxies live, which would followed by successive SOM training until the entire relevant regions of model space have been subsequently mapped using the minimum number of objects. This approach would then serve as a way of interpolating between sparse sampling of the relevant manifold, which could be used to improve our understanding of galaxy properties and accelerate SED fitting techniques. We are currently looking into a number of these possibilities both as they apply to improving photometric redshifts as well as SED fitting more generally.

## ACKNOWLEDGEMENTS

JSS would like to thank Michael Brown, Peter Capak, and Daniel Masters for insightful discussions as well as Charles Alcock for supervising the senior thesis course where a portion of this work was completed. JSS is grateful for financial support from the CREST program, which is funded by the Japan Science and Technology (JST) Agency. This work has benefited extensively from access to Harvard University's Odyssey computing cluster, which is supported by the FAS Division of Science's Research Computing Group.

## APPENDIX A: LINEARIZED FUZZY TEMPLATES

Ignoring rest-frame dust attenuation, the “baseline” model photometry  $\mathbf{F}_{\text{base,gal}}$  for a galaxy at a given redshift can be written as a linear combination of the photometric fluxes derived from the underlying galaxy and emission line templates,

$$(\mathbf{F}_{\text{base,gal}})(z) = \sum_{\text{gal}} c_{\text{gal}} \mathbf{F}_{\text{gal}}(z) + \sum_{\text{lines,gal}} \Delta \text{EW}_{\text{lines}}^{\text{gal}} \mathbf{F}_{\text{lines}}^{\text{gal}}(z). \quad (\text{A1})$$

where

$$\mathbf{F}_{\text{gal}}(z) = \int_{\nu_z} S_{\nu,\text{gal}}(\nu) 10^{-A_{\text{IGM}}(\nu,z)/2.5} [\mathbf{w}(z) \odot \mathbf{T}(\nu)] \nu^{-1} d\nu$$

$$\mathbf{F}_{\text{lines}}^{\text{gal}}(z) = \int_{\nu_z} S_{\nu,\text{lines}}^{\text{gal}}(\nu) 10^{-A_{\text{IGM}}(\nu,z)/2.5} [\mathbf{w}(z) \odot \mathbf{T}(\nu)] \nu^{-1} d\nu$$



are photometric fluxes (including dust attenuation from the IGM) for each galaxy and its corresponding set of emission line templates  $S_{\nu, \text{lines}}^{\text{gal}}$ ,  $\odot$  indicates element-wise multiplication, and  $\mathbf{w}(z) = 1/\int_{\nu_z} \mathbf{T}(\nu)\nu^{-1}d\nu$  is the associated set of inverse filter weights.

As a large set of archetypes should be a representative sample of the corresponding regions of color space, we can ignore linear combinations of galaxy templates, leaving us with the simplified expression

$$s\mathbf{F}_{\text{base,gal}}(z) = s\mathbf{F}_{\text{gal}}(z) + \sum_{\text{lines}} \Delta\text{EW}_{\text{lines}}^{\text{gal}} s\mathbf{F}_{\text{lines}}^{\text{gal}}(z), \quad (\text{A2})$$

where we have included the general scaling factor  $s$  to emphasize that our photometry must be properly scaled before comparing with observed photometry  $\mathbf{F}_{\text{obs}}$ .

Although the impact of attenuation from rest-frame dust screen for a given template is wavelength-dependent, its impact in each band can be approximated as  $\exp[\mathcal{R}_{\text{dust}}(z)\Delta E(B-V)]$ , where  $\mathcal{R}_{\text{dust}}(z)$  is a collection of negative numbers determined by the given dust template  $k_{\text{dust}}(\nu)$  and the corresponding filter set  $\mathbf{T}(\nu)$  derived at a given redshift. As we are only considering the perturbative impact of dust on the few-percent level, we choose to Taylor-expand this expression about 0. Ignoring all  $\mathcal{O}( > 1 )$  terms (including cross-products), we get

$$s\mathbf{F}_{\text{model}}(z) = s\mathbf{F}_{\text{gal}}(z) + \Delta E(B-V) s\mathbf{F}_{\text{dust}}(z) + \sum_{\text{lines}} \Delta\text{EW}_{\text{lines}}^{\text{gal}} s\mathbf{F}_{\text{lines}}^{\text{gal}}(z), \quad (\text{A3})$$

where  $\mathbf{F}_{\text{dust}}(z) = \mathcal{R}_{\text{dust}}(z) \odot \mathbf{F}_{\text{gal}}(z)$  is our new “dust photometry” term.

As our first-order expansion translates dust components that are linear in *magnitudes* to be linear in *flux*, we can use it to probe additional additive dust components that might improve our fitting routine. In particular, we consider a modified dust template of the form

$$k'_{\text{dust}}(x) = k_{\text{dust}}(x) + c_b k_{\text{bump}}(x|x_0, \gamma), \quad (\text{A4})$$

where  $x \equiv 1/\lambda$  is measured in  $\mu\text{m}^{-1}$  and

$$k_{\text{bump}}(x|x_0, \gamma) = \frac{x^2}{(x^2 - x_0^2)^2 + x^2\gamma^2}, \quad (\text{A5})$$

is a Lorentzian-like function where  $x_0$  and  $\gamma$  the central position and width of the Lorentzian feature, respectively. This “bump”, most often located around 2175 Å, is a common observed feature among both Galactic (Fitzpatrick & Massa 2007) and extragalactic (Kriek & Conroy 2013; Scoville et al. 2015) sources. As its amplitude  $c_b$  can vary widely, ignoring its impact on observed photometry can significantly bias quantities derived through SED fitting (Kriek & Conroy 2013).

Inserting our new dust curve into our original expression for  $\mathbf{F}_{\text{model}}$ , we get the final expression

$$s\mathbf{F}_{\text{model}}(z) = s\mathbf{F}_{\text{gal}}(z) + \Delta E(B-V) s\mathbf{F}_{\text{dust}}(z) + \Delta c'_b s\mathbf{F}_{\text{bump}}(z) + \sum_{\text{lines}} \Delta\text{EW}_{\text{lines}}^{\text{gal}} s\mathbf{F}_{\text{lines}}^{\text{gal}}(z), \quad (\text{A6})$$

where  $c'_b = \Delta E(B-V) \times \Delta c_b$  and  $\mathbf{F}_{\text{bump}}(z) = \mathcal{R}_{\text{bump}}(z) \odot \mathbf{F}_{\text{gal}}(z)$ . While this decouples the two dust terms if they are treated independently, it is still an improvement over excluding the feature entirely.

The corresponding log-likelihood when fitting normally distributed data with a series of independent Gaussian priors is equivalent to a modified  $\chi^2$  metric,

$$\chi_{\text{mod}}^2(\boldsymbol{\theta}, \boldsymbol{\phi}, s) \equiv \sum_i \sigma_i^{-2} [\Delta F_i(\boldsymbol{\theta}, \boldsymbol{\phi}, s)]^2 + \sum_j \left( \frac{\phi_j}{\sigma_{\phi_j}(\boldsymbol{\theta})} \right)^2 \quad (\text{A7})$$

where  $\sigma_i^2 = \sigma_{\text{obs},i}^2 + \sigma_{\text{model},i}^2$  is the  $i$ th component of the total variance,<sup>5</sup>  $\Delta F_i(\boldsymbol{\theta}, \boldsymbol{\phi}, s) = F_{\text{obs},i} - sF_{\text{model},i}(\boldsymbol{\theta}, \boldsymbol{\phi})$  is the  $i$ th flux residual,  $\boldsymbol{\theta} = \{z, \text{gal}, \text{dust}\}$  contains the non-linear parameters of interest,  $\boldsymbol{\phi} = \{E(B-V), c'_b, \{\Delta\text{EW}\}_{\text{lines}}\}$  contains the linear nuisance parameters,  $\sigma_{\phi_j}(\boldsymbol{\theta})$  is the standard deviation of the corresponding Gaussian prior for a given  $\boldsymbol{\theta}$ , the sum over  $i$  is taken over all observed bands, and the sum over  $j$  is taken over all relevant nuisance parameters.

For  $\boldsymbol{\theta}$  and  $\boldsymbol{\phi}$  fixed, we can marginalize over  $s$  to minimize  $\chi_{\text{mod}}^2(s|\boldsymbol{\theta}, \boldsymbol{\phi})$ , giving us

$$s = \sum_i \sigma_i^{-2} F_{\text{obs},i} F_{\text{model},i} / \sum_i \sigma_i^{-2} F_{\text{model},i}^2, \quad (\text{A8})$$

which can be calculated prior to computing the actual  $\chi_{\text{mod}}^2$  value.

While marginalizing over  $s$  is a simple one-step process, marginalizing over the linear nuisance parameters  $\boldsymbol{\phi}$  for fixed  $s$  and  $\boldsymbol{\theta}$  is slightly more computationally demanding. We first re-write  $sF_{\text{model},i}(\boldsymbol{\theta}, \boldsymbol{\phi})$  as

$$sF_{\text{model},i}(\boldsymbol{\theta}, \boldsymbol{\phi}) = sF_{\text{gal},i}(\boldsymbol{\theta}) + \sum_j s\mathbf{X}_{ij}(\boldsymbol{\theta})\phi_j, \quad (\text{A9})$$

where

$$\mathbf{X}(\boldsymbol{\theta}) = \begin{pmatrix} F_{1,\text{dust}}(z) & F_{1,\text{bump}}(z) & F_{1,\text{line-1}}^{\text{gal}}(z) & \dots & F_{1,\text{line-m}}^{\text{gal}}(z) \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ F_{n,\text{dust}}(z) & F_{n,\text{bump}}(z) & F_{n,\text{line-1}}^{\text{gal}}(z) & \dots & F_{n,\text{line-m}}^{\text{gal}}(z) \end{pmatrix}$$

is the  $n \times (2+m) = N_{\text{filt}} \times (2+N_{\text{lines}})$  matrix of pre-computed coefficients for a given  $\boldsymbol{\theta}$  and  $F_i$  is the  $i$ th component (i.e. filter) of the corresponding vector.

Minimizing  $\chi_{\text{mod}}^2(\boldsymbol{\phi}|\boldsymbol{\theta}, s)$  involves solving a system of linear equations of the form

$$\left( \mathbf{X}(\boldsymbol{\theta})^T \mathbf{W}_{\text{obs}} \right) \Delta \mathbf{F}_{\text{gal}}(\boldsymbol{\theta}) = \left( \mathbf{X}(\boldsymbol{\theta})^T \mathbf{W}_{\text{obs}} \mathbf{X}(\boldsymbol{\theta}) + \mathbf{W}_{\boldsymbol{\phi}}(\boldsymbol{\theta}) \right) \boldsymbol{\phi}, \quad (\text{A10})$$

where  $\Delta \mathbf{F}_{\text{gal}}(\boldsymbol{\theta}) = \mathbf{F}_{\text{obs}} - s\mathbf{F}_{\text{gal}}(\boldsymbol{\theta})$  is the baseline galaxy flux residual,  $\mathbf{W}_{\text{obs}} = \text{diag}(\dots, \sigma_i^{-2}, \dots)$  is the associated observational weight matrix,  $\mathbf{W}_{\boldsymbol{\phi}} = \text{diag}(\dots, \sigma_{\phi_j}^{-2}, \dots)$  is the prior weight matrix, and  $T$  is the transpose operator. This gives us an simple iterative scheme for minimizing  $\chi_{\text{mod}}^2(\boldsymbol{\theta}, \boldsymbol{\phi}, s)$  with respect to  $s$  and  $\boldsymbol{\phi}$  for a given choice of  $\boldsymbol{\theta}$ .

## References

- Abazajian K. N., et al., 2009, *ApJS*, **182**, 543  
 Albrecht A., et al., 2006, *ArXiv Astrophysics e-prints*,  
 Almosallam I. A., Lindsay S. N., Jarvis M. J., Roberts S. J., 2015, preprint, ([arXiv:1505.05489](https://arxiv.org/abs/1505.05489))

<sup>5</sup> For the remainder of the paper, we will assume  $\sigma_{\text{model}} = 0$  such that  $\sigma = \sigma_{\text{obs}}$ .

Benítez N., 2000, *ApJ*, **536**, 571

Bernton E., Yang S., Chen Y., Shephard N., Liu J. S., 2015, preprint, ([arXiv:1506.08852](#))

Bolzonella M., Miralles J.-M., Pelló R., 2000, *A&A*, **363**, 476

Bonnett C., 2015, *MNRAS*, **449**, 1043

Bordoloi R., et al., 2012, *MNRAS*, **421**, 1671

Brown M. J. I., et al., 2014, *ApJS*, **212**, 18

Calzetti D., Armus L., Bohlin R. C., Kinney A. L., Koornneef J., Storchi-Bergmann T., 2000, *ApJ*, **533**, 682

Carrasco Kind M., Brunner R. J., 2013, *MNRAS*, **432**, 1483

Carrasco Kind M., Brunner R. J., 2014a, *MNRAS*, **438**, 3409

Carrasco Kind M., Brunner R. J., 2014b, *MNRAS*, **441**, 3550

Coleman G. D., Wu C.-C., Weedman D. W., 1980, *ApJS*, **43**, 393

Collister A. A., Lahav O., 2004, *PASP*, **116**, 345

Cool R. J., et al., 2013, *ApJ*, **767**, 118

Cunha C. E., Huterer D., Lin H., Busha M. T., Wechsler R. H., 2014, *MNRAS*, **444**, 129

Dahlen T., et al., 2013, *ApJ*, **775**, 93

Doucet A., De Freitas N., Gordon N., 2001, Sequential monte carlo methods in practice. Springer-Verlag

Elliott J., de Souza R. S., Krone-Martins A., Cameron E., Ishida E. E. O., Hilbe J., 2015, preprint, ([arXiv:1507.01293](#))

Feldmann R., et al., 2006, *MNRAS*, **372**, 565

Fitzpatrick E. L., Massa D., 2007, *ApJ*, **663**, 320

Foreman-Mackey D., Hogg D. W., Lang D., Goodman J., 2013, *PASP*, **125**, 306

Gerdes D. W., Sypniewski A. J., McKay T. A., Hao J., Weis M. R., Wechsler R. H., Busha M. T., 2010, *ApJ*, **715**, 823

Goodman J., Weare J., 2010, *Communications in Applied Mathematics and Computer Science*, **5**, 65

Hildebrandt H., et al., 2010, *A&A*, **523**, A31

Hoyle B., 2015, preprint, ([arXiv:1504.07255](#))

Hoyle B., Rau M. M., Bonnett C., Seitz S., Weller J., 2015, *MNRAS*, **450**, 305

Ilbert O., et al., 2006, *A&A*, **457**, 841

Ilbert O., et al., 2013, *A&A*, **556**, A55

Ivezic Z., et al., 2008, preprint, ([arXiv:0805.2366](#))

Johnson S. P., Wilson G. W., Tang Y., Scott K. S., 2013, *MNRAS*, **436**, 2535

Kennicutt Jr. R. C., 1998, *ARA&A*, **36**, 189

Kennicutt R. C., Evans N. J., 2012, *ARA&A*, **50**, 531

Kinney A. L., Calzetti D., Bohlin R. C., McQuade K., Storchi-Bergmann T., Schmitt H. R., 1996, *ApJ*, **467**, 38

Kohonen T., 1982, *Biological Cybernetics*, **43**, 59

Kohonen T., 2001, Self-Organizing Maps. Springer, 2001, xx, 501 p. Springer series in information sciences

Kriek M., Conroy C., 2013, *ApJ*, **775**, L16

Laureijs R., et al., 2011, preprint, ([arXiv:1110.3193](#))

Li I. H., Yee H. K. C., 2008, *AJ*, **135**, 809

Madau P., 1995, *ApJ*, **441**, 18

Masters D., et al., 2015, preprint, ([arXiv:1509.03318](#))

Ménard B., Scranton R., Schmidt S., Morrison C., Jeong D., Budavari T., Rahman M., 2013, preprint, ([arXiv:1303.4722](#))

Newman J. A., et al., 2015, *Astroparticle Physics*, **63**, 81

Polletta M., et al., 2007, *ApJ*, **663**, 81

Prevot M. L., Lequeux J., Prevot L., Maurice E., Rocca-Volmerange B., 1984, *A&A*, **132**, 389

Sadeh I., Abdalla F. B., Lahav O., 2015, preprint, ([arXiv:1507.00490](#))

Sánchez C., et al., 2014, *MNRAS*, **445**, 1482

Scoville N., et al., 2007, *ApJS*, **172**, 1

Scoville N., Faisst A., Capak P., Kakazu Y., Li G., Steinhardt C., 2015, *ApJ*, **800**, 108

Speagle J. S., Capak P. L., Eisenstein D. J., Masters D. C., Steinhardt C. L., 2015, preprint, ([arXiv:1508.02484](#))

Tomczak A. R., et al., 2014, *ApJ*, **783**, 85

Wittman D., 2009, *ApJ*, **700**, L174

This paper has been typeset from a  $\text{\TeX}/\text{\LaTeX}$  file prepared by the author.